

DELL

MAY 2003 • SPECIAL ISSUE

POWER SOLUTIONS

THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS




Migrating to Microsoft Windows Server 2003

Inside This Issue:

Exploring Microsoft .NET

Clustering for high availability

Examining PowerEdge performance



Lower Cost.
Better Performance.
You do the math.

IT customers are saving more.

IT Managers are shaving significant costs off their systems, improving reliability and protecting their infrastructure investment. What's their secret?

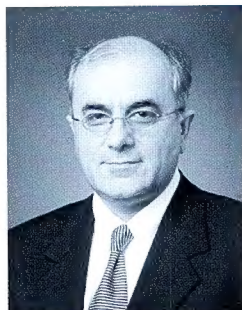
15K disc drives. Today's 15K drive achieves a lower total cost of ownership than yesterday's 10K drive. Discover how the 15K:

- Decreases space requirements
- Lowers annual power and cooling costs
- Reduces acquisition and administration costs
- Preserves your investment longer

Decrease cost. Increase reliability. The proof is in the data.

Download the Seagate *Economies of Speed* white paper, the Dell *Power Solutions* article with Dell™ test results, and gain insight into third-party test results—all located at specials.seagate.com/15K.





Dell and Microsoft:

A Partnership Working for You

Businesses demand technologies that will enable them to operate more efficiently. Selecting the software most likely to streamline business processes and resources is a challenge for IT organizations—and deploying it with the most effective hardware available is often just as critical. The decision to adopt a preferred software standard often depends, in part, on its interoperability with the hardware on which it will be installed.

Dell and Microsoft understand that failure is not an option for businesses deploying new IT infrastructures or upgrading existing implementations. That's why the release of the Microsoft® Windows® Server 2003 operating system builds on the strengths of Windows 2000 Server to provide enterprise computing environments with high levels of manageability, reliability, scalability, and security.

The new operating system offers improved support for Microsoft Active Directory® services, eases the management of server and storage resources in Microsoft .NET infrastructures, leverages the reliability of Windows 2000 Server, and supports the development of scalable Extensible Markup Language (XML) Web services. Windows Server 2003 is designed to be the most secure Microsoft operating system ever shipped.

Holistic product release ensures interoperability

Perhaps more importantly, Dell and Microsoft have collaborated extensively to pull together a complete, proven IT solution. During the Microsoft development of Windows Server 2003, Dell and Microsoft worked together closely at technical and program management levels to ensure full compatibility of Windows Server 2003 on Dell™ PowerEdge™ servers and Dell storage products. This cooperation included:

- Comprehensive testing of beta operating system releases at Dell and its partners
- Incorporating Dell hardware into the test beds at Microsoft

- Testing prototype and production-level PowerEdge servers and hardware in Dell and Microsoft laboratories
- Working with early adopters through the Microsoft Joint Development Program (JDP)

To coincide with the final release of Microsoft Windows Server 2003, Dell is releasing the new operating system on all factory-shipped Dell PowerEdge servers, and extending support to many previous-generation servers. Dell support for Windows Server 2003 also includes peripherals, the Dell OpenManage™ suite of management tools, and the Dell | EMC line of Fibre Channel storage area network (SAN) products.

The extensive testing and early-adopter feedback that went into this holistic release will help ensure the smooth interoperability of Windows Server 2003 with Dell products. Thanks to the close collaboration between Dell and Microsoft, IT departments will be able to move forward with confidence—assured that they have selected products designed from the ground up to operate efficiently together, and that they have the knowledge to deploy them. By supporting the development of Microsoft Windows Server 2003, Dell is enabling today's enterprises to build the dependable infrastructures needed to meet the business challenges of tomorrow.

Reza Rooholamini, Ph.D.
Director of Enterprise Solutions Engineering
Dell Product Group

EDITORIAL

EDITOR-IN-CHIEF | Eddie Ho

MANAGING EDITOR | Lea Anne Bantsari

CONTRIBUTING EDITORS | Stephanie Dobler, Liza Graffeo,
Debra McDonald, Vicki Van Ausdall

ART

ART DIRECTOR | Mark Mastroianni

DESIGNERS | Glen Abrahams, Phu Tran, Amy Vest, Cynthia Webb

COVER DESIGN | Glen Abrahams

ONLINE

WEB PRODUCTION MANAGER | Brad Klenzendorf

SUBSCRIPTION SERVICES

SUBSCRIPTION SERVICES | Lisa Ho

Subscriptions are free to qualified readers who complete the subscription card found in each issue. To subscribe or change your address, complete and return the business reply card in this issue to Dell Computer Corporation, c/o Lisa Ho, One Dell Way, RR #5, Box 8503, Round Rock, TX, 78682, USA, or visit us at www.dell.com/powersolutions_subscribe.

About Dell

Dell Computer Corporation, headquartered in Round Rock, Texas, near Austin, is the world's leading direct computer systems company. Dell is one of the fastest growing among all major computer systems companies worldwide, with approximately 40,000 employees around the globe. Dell uses the direct business model to sell its high-performance computer systems, workstations, and storage products to all types of enterprises. For more information, please visit our Web site at www.dell.com.

Dell™, OpenManage™, PowerConnect™, PowerEdge™, PowerVault™—Dell Computer Corporation; 3Com™—3Com Corporation; Adaptec™—Adaptec, Inc.; Altec™—Altec Web Systems, Inc.; DLT™—Benchmark Storage Innovations, Inc.; Broadcom™, NetXtreme™—Broadcom Corporation; CLARiON™, Connectrix™, ControlCenter™, EMC™, PowerPath™, Symmetrix™—EMC Corporation; cLAN™, Emulex™, Giganet™—Emulex Corporation; Extreme Networks™, Summit™—Extreme Networks Corporation; Fujitsu™—Fujitsu Limited Corporation; IEEE™—Institute of Electrical and Electronics Engineers, Inc.; Intel™, Itanium™, Pentium™, Xeon™—Intel Corporation; DB2™, Domino™, IBM™, Lotus™—International Business Machines Corporation; LSI Logic™—LSI Logic Corporation; Active Directory™, ActiveX™, JScript™, Microsoft™, Visual SourceSafe™, Visual Studio™, Windows™, Windows NT™—Microsoft Corporation; NetWare™, Novell™—Novell, Inc.; UNIX™—The Open Group; Oracle™—Oracle Corporation; QLogic™—QLogic Corporation; Rosetta Genomics™—Rosetta Genomics Ltd.; Seagate™—Seagate Technology, Inc.; Linear Tape-Open™, LTO™—Seagate Technology, Inc.; International Business Machines Corporation, and Hewlett-Packard Company; VERITAS Backup Exec™—VERITAS Software Corporation; TapeWare™, Yosemite Technologies™—Yosemite Technologies, Inc.; NetBench™, WebBench™—Ziff Davis Media Inc.; Ziff Davis™—Ziff Davis Publishing Holdings Inc. Other company, product, and service names may be trademarks or service marks of others.

Dell Power Solutions is published quarterly by the Dell Product Group, Dell Computer Corporation, One Dell Way, Round Rock, Texas 78682. This publication is also available online at www.dell.com/powersolutions. No part of this publication may be reprinted or otherwise reproduced without permission from the editor. Dell does not provide any warranty as to the accuracy of any information provided through *Dell Power Solutions*. Opinions expressed in this magazine may not be those of Dell. The information in this publication is subject to change without notice. Any reliance by the end user on the information contained herein is at the end user's risk. Dell will not be liable for information in any way, including but not limited to its accuracy or completeness. Dell does not accept responsibility for the advertising content of the magazine nor for any claims, actions, or losses arising therefrom. Goods, services, and/or advertisements within this publication other than those of Dell are not endorsed by or in any way connected with Dell Computer Corporation.

Copyright © 2003 Dell Computer Corporation. All rights reserved. Printed in the U.S.A.

Special Issue | May 2003



SPECIAL ISSUE | MAY 2003

EXECUTIVE NOTES

1 | Dell and Microsoft: A Partnership Working for You By Reza Rooholamini, Ph.D.

EDITOR'S COMMENTS

4 | Microsoft .NET: A New Framework for Enterprise Agility By Eddie Ho

OS ENVIRONMENT

6 | Exploring Microsoft .NET

By Michael P. Kennedy

10 | Keys to Successful Deployment of .NET Applications

By Scott Jamison

16 | Evaluating Windows Server 2003 Migration Options

By Mike Owens

21 | Dell and Microsoft: Partnering to Deliver Windows Server 2003 on PowerEdge Servers

By Lara Benson and Keith Matteson

25 Guide to Deploying Microsoft Windows Server 2003 on Dell PowerEdge Servers

By the Dell Server Operating Systems Engineering Group

33 Deploying Dell OpenManage on Microsoft Windows Server 2003

By Paul Rad; Mohammad Dhedhi; Joseph Ondrejik, Ph.D.; and Candace Hobson

38 Using Windows System Resource Manager for Server Consolidation

By Ranjith Purush, Chris Stackhouse, and Paul Rad

43 Examining Windows Server 2003 Performance on the Dell PowerEdge 4600 Server

By Quynh O. Nguyen and Min-John Lee

47 Using PCI Hot Plug on Dell PowerEdge Servers Running Windows Server 2003

By Scott M. Callaway and Faisal Ahmed

56 Using Microsoft Cluster Service to Enhance Dell Clusters

By Mahmoud Ahmadian and Daniel Moges

59 Microsoft SQL Server 2000 Scalability: The Rosetta Genomics 10 TB Human Genome Database

By Joe Belson and Hao Chen

64 Sizing Microsoft Exchange 2000 on Dell PowerEdge Servers and Dell | EMC Fibre Channel Storage Systems

By Guy Westbrook

APPLICATION ENVIRONMENT

51 High-Availability Clustering Capabilities in Windows Server 2003

By Mike Kosacek, Nam Nguyen, and Mark Tibbs

Advertiser Index

Dell	5, 9, 18-19, 23, 37
Microsoft Corporation	12-13, inside back cover, back cover
Seagate Technology LLC	inside front cover



Microsoft .NET:

A New Framework for Enterprise Agility

Technology providers, industry strategists, analysts, and users seldom agree on what constitutes the best IT solutions, but the idea of Web services seems to have united all parties. The capabilities of the Internet have changed business transactions and collaboration, making Web-based intersystem communication a necessity for many enterprises. With the acceleration of business and IT integration, both within the enterprise and between enterprises, a primary challenge is to create a standardized platform for system-level collaboration.

However, most computing environments today are autonomous and permit only limited collaboration with other environments. The Microsoft® .NET Framework is designed to encourage a standardized interface, allowing computing environments to communicate over the Internet. It also offers a rich set of application programming interfaces (APIs) that forms the basis of the Microsoft Windows® Server 2003 operating system.

Windows Server 2003 offers excellent reliability, availability, scalability, and security for enterprise computing environments. The operating system is designed to ease enterprise deployment and to enhance systems management.

This special issue of *Dell Power Solutions* focuses on the Microsoft .NET Framework and the Windows Server 2003 operating system. Articles examine migration to and deployment of these Microsoft technologies in Dell™ PowerEdge™ server environments. Topics include the following:

- Overview of .NET
- General migration strategy for Windows Server 2003 on Dell PowerEdge servers
- Techniques for managing a Windows Server 2003 environment
- Clustering servers for high availability
- Performance measures for Windows Server 2003
- Deployment considerations for Microsoft Exchange

We hope this special issue of *Dell Power Solutions* helps jump-start your Windows Server 2003 migration.

A poster included in the magazine details cluster deployment scenarios for Dell systems. In addition, the May issue of our business IT magazine, *Dell Insight*, takes an industry perspective on the business agility of .NET. For more information about this publication, please visit <http://www.dell.com/dellinsight>.

In conjunction with Microsoft's launch of Windows Server 2003, Dell will host a series of Webcasts and seminars discussing the Microsoft .NET platform. The first Webcast, "Microsoft .NET—The Next-Generation OS," begins Wednesday, May 14, at 12 P.M. (CST), followed by a different Webcast each Wednesday at noon for the next five weeks. The six sessions will focus on the following topics: basic understanding of Windows Server 2003, business advantages and deployment scenarios for Windows Server 2003, Microsoft Exchange 2003, Microsoft SQL Server 2003, the Microsoft .NET programming environment, and migration and server consolidation. For more information about Dell Webcasts, visit <http://www.dell.com/powersolutions> and select "Webcasts."

In June, the *Dell Power Solutions* seminar series will launch a half-day seminar entitled "Building Better Business Agility with Microsoft .NET Solutions." This seminar opens in Asia and will take place in various countries around the world during the following six months. For more information about the .NET seminars, visit <http://www.dell.com/powersolutions> and select "Seminars."

Dell remains committed to providing timely, practical information that enhances computing infrastructures. We appreciate your feedback. Together we can improve the quality of IT.

Eddie Ho
Editor-in-Chief
www.dell.com/powersolutions

Dell High Performance Computing Clusters.

Want to reinvent your data center?
Follow the leader.



Dell servers use cost effective, industry-standard technology such as Intel® Xeon™ processors.

Dell | Enterprise

The world of supercomputing has changed. No longer does a business have to rely on expensive, proprietary systems to run mainframe applications. Let Dell, the leader in high performance computing clusters, bring a singular focus on TCO to your data center. Built on flexible, cost-effective Intel® Xeon™ processor technology, Dell High Performance Computing Clusters can help you make an impact within your organization that goes far beyond IT. What can Dell High Performance Computing Clusters provide your business?

- **Instant scalability.** Dell HPCCs give you the ability to grow in overall capacity and to meet high usage demand as the need arises.
- **Enhanced availability.** With a Dell HPCC, the system as a whole stays highly available, even with the removal of any single failure point in hardware or software.
- **Enterprise-level service and support.** Full, 24/7 commitment from design to implementation to on-site¹ or online support.
- **Improved system manageability.** Dell's Intelligent IT portfolio gives you the tools and services to automate the deployment, management and maintenance of your IT infrastructure. The result? Lower TCO.

For nearly 20 years, we've revolutionized the way the world buys and manages technology. Now see how we're revolutionizing the world of High Performance Computing. To put a Dell HPCC to work for your business, visit dell.com/hpcc.



High performance computing that can lower TCO. Easy as **DELL**

Visit www.dell.com/hpcc

¹Service may be provided by third-party. Technician will be dispatched if necessary following phone-based troubleshooting. Subject to parts availability, geographical restrictions and terms of service contract. Service timing dependent upon time of day call placed to Dell. U.S. only. Intel, the Intel logo and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Dell, the Dell logo and PowerEdge are registered trademarks of the Dell Computer Corporation. ©2003 Dell Computer Corporation. All rights reserved.

Exploring Microsoft .NET

The Microsoft® .NET architecture, slated to become the basis of Microsoft applications, tools, and servers, offers many advantages to developers and users: .NET helps to facilitate interoperability, ease system integration, and enhance communication between systems and business partners. Although .NET has aroused much interest, some confusion still surrounds its function and purpose. This article explains what .NET is and what it can do.

BY MICHAEL P. KENNEDY

Microsoft® .NET is software for connecting information, people, systems, and devices. The Microsoft .NET architecture includes a development and run-time environment that enables developers to quickly build applications that interoperate across languages, platforms, and devices. For IT managers, .NET eases integration of new and existing systems across the organization and externally with partners. Consumers using smart devices benefit from the .NET cross-platform capability by gaining greater access to information.

Ensuring interoperability with XML

.NET employs Extensible Markup Language (XML), a standard ratified by the World Wide Web Consortium (W3C). W3C is the standards organization responsible for developing common protocols that promote the Web's evolution and ensure its interoperability. Microsoft and other computer industry leaders work with W3C to promote the adoption of new features and standards.

Using markup tags to add information to data

Before XML, when two applications needed to communicate with each other, they would use a proprietary communication method. Developers would write one computer function to send data and another to receive data. The data was separated into different sections—called fields—by commas, tabs, or some other indicator. A data stream in a typical interface might look like this:

```
"ABC47-Z", "100", "STL", "C", "3", "28"
```

This method worked well, but if another application also wanted to communicate, developers needed to write another send function and receive function to accommodate that application's proprietary interface and language.

In XML, the same data stream shown in the previous paragraph is represented quite differently:

```
<InventoryItem>
  <PartNum>ABC47-Z</PartNum>
  <Quantity>100</Quantity>
  <Warehouse>STL</Warehouse>
  <Zone>C</Zone>
  <Aisle>3</Aisle>
  <Bin>28</Bin>
</InventoryItem>
```

Now the meaning of each piece of data can be inferred, because the markup tags that accompany it indicate what the data represents. An application that understands XML uses the tags to easily locate and process only the data it has been directed to find, and ignores the rest.

The difference between the two methods of handling data has dramatic implications for the applications that exist on either end of a data transmission. In the former example, no clues indicate what the data represents, and all the rules for finding and processing this data must be

resident in the applications on either end. Almost without exception, if the structure of the data changes, the applications on either end must change.

Interpreting data with a parser

Applications using XML can rely on a standard XML interpreter, called a parser, to get the right data out of the stream. If an application has an XML parser and needs the `PartNum` and `Quantity` data, for example, it will retrieve just those fields. If the data structure is changed, perhaps to include extra fields, the application will not need to change; the XML parser will still be able to provide the data to the application.

The use of XML results in more flexible applications, which can reduce application maintenance costs. The savings can be substantial, because application maintenance constitutes a significant portion of IT expenses in many organizations.

Facilitating communication between applications through XML Web services

XML can be used to mark up data, but applications still need a way to communicate with each other. XML Web services provide this ability. These services constitute reusable pieces of software that interact over the network through industry-recognized standards. XML Web services are made available from a business's Web server for Web users or other Web-connected programs. HTTP, a globally adopted standard used in intranets, extranets, and the Internet, is the protocol used by XML Web services to allow applications to communicate with each other.

Delivering data using the SOAP protocol

Simple Object Access Protocol (SOAP) is used to facilitate the delivery of XML data using HTTP. This protocol allows applications to call each other in a standard, loosely coupled manner that facilitates building and distributing applications across the Internet or an intranet. SOAP can be thought of as an envelope for the data. The data goes inside the SOAP envelope and SOAP guides the data from one application to the other using HTTP. SOAP knows the address of the application that will receive the XML data, process the data, and return the requested information to the originating application.

Publishing, finding, and using Web services through UDDI

Universal Description, Discovery, and Integration (UDDI), an XML-based distributed directory similar to a phone book, enables businesses to publish, find, and use Web services. Companies can discover technical details about working with other Web services, and post details so that other companies can use theirs. Developers can use UDDI to keep track of Web services for use inside the company. The ability to develop a Web service to an

application once and reuse it again and again for other applications saves a significant amount of time.

Simplifying application integration using Web services

Companies can now create XML Web services as the interface to their applications, and these applications can then communicate more effectively with each other. As a result, users can access critical data from multiple sources running in different applications, on different hardware, and from different companies. Companies can transfer information using industry standards. These standards simplify data exchange, making partnering more efficient and cost-effective. Web services also simplify application integration, reducing support costs.

Integrating support for Web services in the enterprise

More than 50 software and technology companies participate in the Web Services Interoperability Organization (WS-I). The WS-I promotes interoperability among XML Web services based on common industry-accepted definitions and related XML standards support. Instead of deciding between competing proprietary technologies, companies can now choose between multiple providers that use an agreed-upon standard. Reliability, total cost of ownership (TCO), and time to implement are important decision criteria when selecting a technology provider.

The .NET platform: Supporting enterprise applications

As businesses extend their computing networks to partners and customers, they need an infrastructure that can readily respond to changing markets while supplying the same reliability, security, and scalability as any traditional enterprise application. The Microsoft .NET platform includes development tools and server infrastructures that take full advantage of XML Web services to support smart devices and enterprise applications.

The Microsoft .NET
platform includes
development tools and
server infrastructures that
take full advantage of
XML Web services to
support smart devices and
enterprise applications.

Building applications using .NET development tools

Developers can use the Microsoft Visual Studio® .NET development system and the .NET Framework to build applications and services on the XML Web services platform. These two tools provide a unified environment for developing, deploying, and running powerful XML Web services that allow for maximum interoperability and minimal integration costs.

Each application and service created using the .NET development tools is a potential building

block for future applications and services. By using the .NET Framework, developers take advantage of open standards, particularly XML, so they do not need to waste time reinventing the wheel.

Visual Studio .NET. Microsoft has reengineered Visual Studio .NET to help developers more easily write distributed applications as XML Web services. Visual Studio .NET provides a shared integrated development environment (IDE) for all the languages within it, helping developers to build applications faster using their languages of choice. Additional enhancements include improved window management, design tools for XML, Visual Studio macros to automate repetitive tasks, and database tools for easier integration of databases.

Using Visual Studio .NET Enterprise Edition, developers can choose from more than 24 programming languages to build applications that integrate data and business logic from existing systems, including SQL Server, Oracle®, and IBM® DB2® databases. This multiplatform support enables companies to extend the useful life of previous investments in technology and programmer skill sets.

Developers can create an end-to-end view of applications while in development, including components and business process flows. Functional and stress-testing tools built into the development environment ensure that applications will model expected performance and scale to meet demand.

The .NET Framework. The Microsoft .NET Framework provides the programming model for building, deploying, and running XML Web services and applications. The application installation used by the .NET Framework solves deployment and operation challenges presented by enterprise-scale applications. In addition, different versions of applications developed using .NET can coexist side by side, eliminating one major cause of application incompatibility.

The common language runtime (CLR) component of the .NET Framework provides a standards-based, multilanguage environment for integrating existing applications with next-generation .NET applications and services. The .NET Framework also contains a hierarchical set of unified class libraries. These libraries include ASP.NET, a component-based version of Active Server Pages (ASP); ADO.NET, a loosely coupled data access subsystem; and Windows® Forms, an environment for building rich Windows applications. These features and others help decrease development complexity, increase the efficiency of developers, and increase the stability and scalability of Microsoft-based business solutions.

Delivering .NET experiences through smart devices

Smart devices include PCs, laptops, workstations, phones, handheld computers, Tablet PCs, and game consoles that can use XML, SOAP, and UDDI to access XML Web services, thereby enabling users to interact with data regardless of their locations. Smart devices remember device usage information, including user preferences and profile information, to customize the user experience and provide .NET experiences. Smart devices also are sensitive to network conditions


and can facilitate seamless user experiences both online and offline. Additionally, a smart device can enable the input method, connection, and data presentation most appropriate for that device.

Maximizing scalable server infrastructures

Scalable servers that deeply integrate XML will help establish the infrastructure for enterprise computing in which software is delivered as a service; where information is accessible by any device, at any time, from any place; and where solutions are fully scalable, programmable, and customizable to provide maximum business value. Microsoft Windows 2000 Server and the .NET Enterprise Server products provide a foundation for future XML Web services. The Windows 2000 Server family features built-in support for XML, the key building block for .NET technology. The .NET Enterprise Server products form a reliable platform for demanding enterprise applications and are well suited for deploying, operating, and managing XML Web services. They also provide broad support for existing and emerging Internet standards.

The upcoming Windows Server 2003 operating system will increase the scalability, reliability, and manageability of the Microsoft server platform. Using the new Secure by Default paradigm, Windows Server 2003 is designed to be the most secure Microsoft operating system ever shipped. The operating system incorporates the .NET Framework, providing inherent support for XML Web services.

Attaining interoperability with .NET

.NET-connected software—including developer tools, applications, and server and client operating systems—provides easier, more integrated communication. True interoperability across platforms, programming languages, and devices offers the opportunity for enhanced interaction and collaboration. As systems become easily linked, the time, effort, and expense of working with partners will decrease. Microsoft .NET helps businesses better connect with their customers, link systems, and provide improved, more convenient access to information. 

Michael P. Kennedy (mkennedy@microsoft.com) is a technology specialist at the Microsoft Technology Center in Austin. He specializes in helping companies understand new technologies from Microsoft and apply new technologies to solve business challenges. Michael has a degree in business from Texas A&M University.

FOR MORE INFORMATION

Microsoft .NET: <http://www.microsoft.com/net>

Microsoft .NET Enterprise servers: <http://www.microsoft.com/servers/>

Microsoft Technology Centers (MTC): <http://www.microsoft.com/usa/mtc>

UDDI: <http://uddi.microsoft.com>

Web Services Interoperability Organization (WS-I): <http://www.ws-i.org>

World Wide Web Consortium (W3C): <http://www.w3.org>

Introducing Dell Blade Servers.

It's amazing what your business can accomplish when you think small.



Dell Enterprise

PowerEdge™ 1655MC Blade Server

Cost-Effective Consolidation Solution

- Up to 2 Intel® Pentium® III Processors at 1.26GHz Per Blade
- Up to 84 1655MC Blades Fit into One Standard 42U Rack
- 128MB - 2GB 133MHz ECC SDRAM
- Supports up to 2 High-Performance SCSI Hard Drives with PERC4/iM Integrated RAID
- Hot Plug Redundant Power and Cooling Standard
- Integrated Management Module and Keyboard, Video and Mouse Switch
- Integrated Layer 2 Managed Ethernet Switches

Consolidate with Dell Blade Servers and dramatically lower TCO. Dell PowerEdge™ 1655MC Blade Servers, powered by Intel® Pentium® III processors, represent the future of server design as well as a more cost-effective way to start consolidating your web, network infrastructure and application servers. Our small, 3U/6 blade chassis is cutting-edge, and it easily fits into your existing power rack and administrative infrastructure. With Dell PowerEdge MC Blade Servers, you can expect:

- **Modular Flexibility:** Consolidate multiple applications in a fraction of the space required by traditional rack designs.
- **Superior ROI:** The 1655MC's smaller, modular design can help improve your ROI when buying three or more. And they're easy to install, maintain and service.
- **Infinite Scalability:** Buy only the number of blade servers you need, without the investment and overhead of much larger, more expensive chassis designs. As your needs grow, simply plug in new blade servers.

So go to www.dell.com/blades and discover the easier way to get on the consolidation bandwagon with Dell PowerEdge MC Blade Servers.



Smarter consolidation solutions. Easy as **DELL**

Visit www.dell.com/blades

Keys to Successful Deployment of .NET Applications

Successful roll-outs of enterprise applications developed using the Microsoft® .NET platform require careful planning, management, and implementation. This article describes common scenarios, controlled environments, and proven techniques to facilitate a smooth application deployment.

BY SCOTT JAMISON

The success of enterprise applications developed using the Microsoft® .NET platform hinges upon a well-designed data center and a meticulously planned deployment process. This article introduces techniques for helping to ensure successful roll-outs of applications.

Deciding on a reference architecture

Application developers usually receive an explicit set of functional requirements for the software. However, several implicit requirements should not be overlooked:

- The new application should be deployed in a secure and controlled manner.
- Web services and Web applications must be consistently available.
- Web services and Web applications must be scalable as business requirements increase.

By following a standard data center configuration, organizations can quickly and efficiently build Internet applications that meet long-term business goals.

A typical enterprise Web application based on .NET consists of three logical tiers:

- **Web application tier:** Provides HTML directly to users so they can interact with the application
- **Web services tier:** Provides business logic and

functional services through .NET remoting or Web services; also known as the business logic tier

- **Data tier:** Provides data storage and retrieval services, typically over Remote Procedure Call (RPC)

When rolling out an application, administrators can decide to implement it with one to three, or more, physical tiers. The reference architecture scenario described in this article uses one physical tier for each logical tier, for a total of three physical tiers. The scenario assumes that usage requirements call for a farm of four Web servers on the Web application tier, a farm of three business logic servers on the Web services tier, and a database cluster on the data tier. In addition to the servers in the data center, the system includes appropriate switching and firewall hardware to accommodate multiple virtual LANs (VLANs) and system security. Figure 1 highlights the key servers and equipment that constitute a typical system architecture for .NET applications.

An appropriate data center design is only the first step in a successful deployment. Key aspects of a .NET enterprise application roll-out include careful attention to testing, staging, distribution, and upgrades.

Planning the deployment process

The system development life cycle (SDLC) includes common activities such as designing, coding, and testing,

as well as activities sometimes given less attention, such as staging, distribution, and in-place upgrades. A well-thought-out deployment plan is critical, yet many development teams treat it as an afterthought. The following section describes what can happen with inadequate planning.

Lack of planning: Unexpected problems

In a common scenario, a development team uses the Microsoft Visual Studio® .NET development platform and achieves successful operation of the application in a test environment. On the day that the system is supposed to go live, the team uses xcopy deployment to move the compiled components and Web pages to their respective servers. However, settings and dependencies that were taken for granted on the development machines, and some Component Object Model (COM) dependencies for legacy support, are suddenly missing, because they were installed and configured with the developer tools.

To solve the problem, the development team hurriedly tries several solutions. The team changes some settings on the production server (but does not record which ones), copies over additional files (again, without noting the details), and finally installs Visual Studio

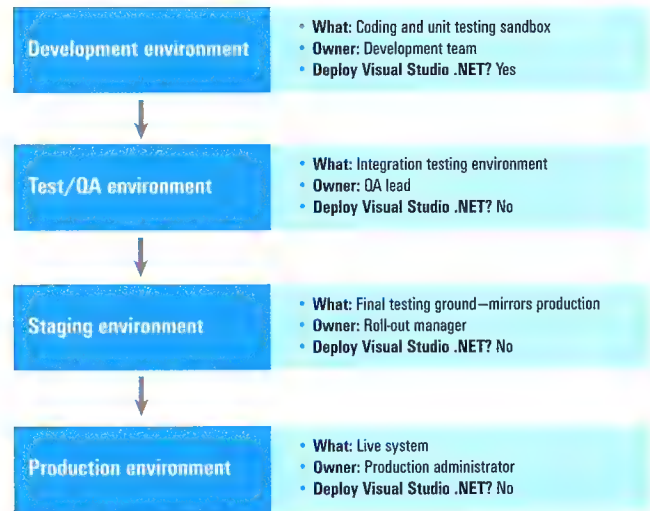


Figure 2. Key environments in a .NET application deployment process

.NET on the production server to make the application work. However, the system now has many problems: developers have access rights to the production environment, creating an internal security and stability risk; development tools and source code reside on the production servers, opening security holes; and a non-repeatable installation procedure now exists, preventing reliable upgrades.

The team in this scenario learns from its mistakes, implements a staging environment, and rolls out the application to the staging environment where it passes all tests. When moved to production, however, the application fails again because the staging servers were not set up to mirror the production environment's load balancing and authentication techniques.

Appropriate planning: A controlled environment

Complex applications such as .NET-compatible enterprise solutions must be tested in a controlled environment separate from the development environment. Deployments should be staged by a roll-out manager, typically on a platform that is identical to, but separate from, the production environment. At minimum, four environments should be used (see Figure 2).

Development environment. Developers perform all coding and unit testing in a development environment. During unit testing, the developer ensures that individual components meet specifications and are operating correctly. Often called a developer's sandbox, this environment usually includes dozens of tools in addition to Visual Studio .NET. Because they lack these tools, the final production servers will not share the same configuration as the development environment.

Test/QA environment. The test environment enables the development and quality assurance teams to test the features of

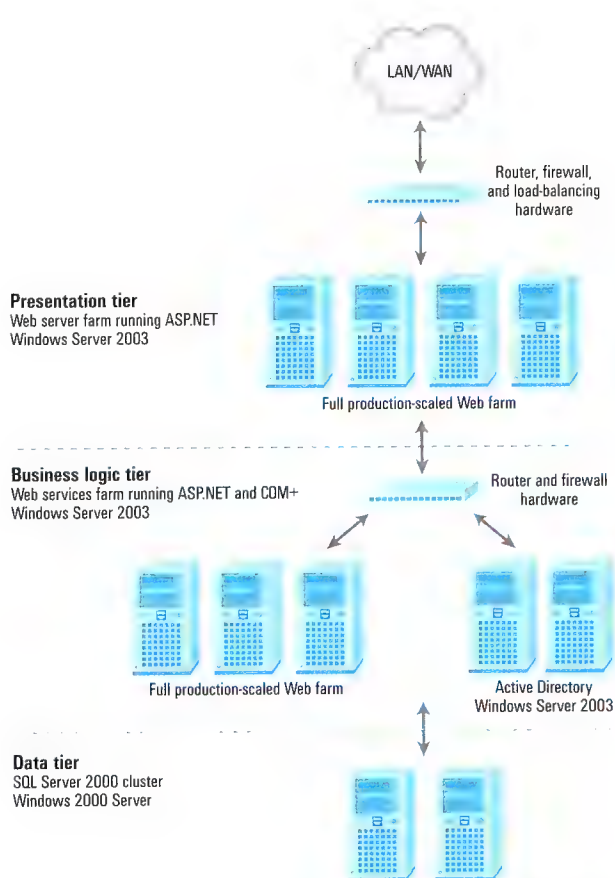


Figure 1. Typical three-tier data center configuration for .NET-developed applications

© 2003 Microsoft Corporation. All rights reserved. Microsoft, the Microsoft Dynamics logo, and the Microsoft Dynamics logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other marks contained herein are the trademarks of their respective owners.

Larger projects, higher goals, greater responsibility.

Introducing Microsoft Windows Server 2003. Do more with less.

You're being asked to do more. You're being asked to do it with less. Microsoft® Windows® Server 2003 is designed to help you manage these opposing forces and deliver powerful software solutions with less time, money, and hassle. For more information, and to get your free evaluation copy of Windows Server 2003 by July 31, 2003, go to microsoft.com/windowsserver2003. **Software for the Agile Business.**

▶ **The London Stock Exchange**, with the help of Accenture, selected Windows Server 2003 as the foundation for their real-time, business-critical, market information delivery system. Using Visual Studio® .NET and the Microsoft .NET Framework, this innovative new system was developed in less than eight months from conception to production and now serves 100,000 terminals worldwide with up to 3,000 transactions per second.

Microsoft

Fewer resources, tighter timelines, shrinking budgets.



Microsoft
Windows Server 2003

the integrated application. Roll-out to the test environment should include obtaining the latest source code from the source control server (typically Microsoft Visual SourceSafe®), packaging the installation, and running the installation program. The team should avoid installing development tools on the test servers, because these tools may introduce dependencies that could mask problems that may surface later.

Staging environment. Often the most neglected of the environments, the staging servers are the most important to a smooth roll-out. The staging environment provides a place to house the application after it has been fully tested in the test environment, but before rolling it out to users. The staging environment provides the last opportunity to find problems before end users see the application. For this reason, the staging environment must be virtually identical to the production environment and use the same operating system, applications, and drivers. The staging environment also should mirror the production network topology and use its lock-down procedures and policies. If the production environment uses clustering or load balancing, these techniques should be used in the staging environment first.

Production environment. The production environment is the live site, the actual network to which end users connect, so downtime must be minimized and changes to applications must proceed smoothly. Even minor changes can cause site failure.

Creating the ideal staging environment

The staging environment should resemble the production environment as closely as possible. Creating the ideal staging environment requires starting with a clean set of servers, the same type already used in production. The staging network should include the same type and number of switches, routers, firewalls, and storage hardware as well. Time and money saved by preventing deployment problems should eventually compensate for the cost of extra equipment.

Although every aspect of production should be mirrored, deployment teams can scale down the staging environment. If four Web servers are in production, two Web servers might be sufficient for the staging environment, enabling proper testing of load balancing and failover scenarios.

The staging environment should include redundant machines; otherwise, certain situations may not be replicated. For example, a production environment could include a stateless Web service that updates a multi-master, synchronized data store, such as the Microsoft Active Directory® directory service. A subsequent read by the Web service could have a problem, because another Active

The staging environment
should resemble the
production environment
as closely as possible.

Directory server—one that has not yet received the data sync—may respond to the request, causing a read failure. Unless the deployment team mimics the production environment with redundant hardware, the team is unlikely to discover this problem. Figure 3 shows a staging environment that suitably mimics the production architecture in Figure 1.

Setup of the staging servers begins with a clean installation of the operating system. Then, during the installation procedure of the application, the deployment team should clearly document every step and action, no matter how small. The team should note the location of application files and all configuration settings, as well as all temporary files, registry settings, and user accounts that need to be created. The completed document should mimic every aspect of the installation process in painstaking detail, with a final section describing steps to uninstall the application.

After completing the installation process, the team should test to ensure that the application is working properly and that the installation document is complete. Then the team should uninstall all software in the staging environment, leaving a bare machine. Next, the team should install the application again, precisely following the installation document. Any action that is not reflected

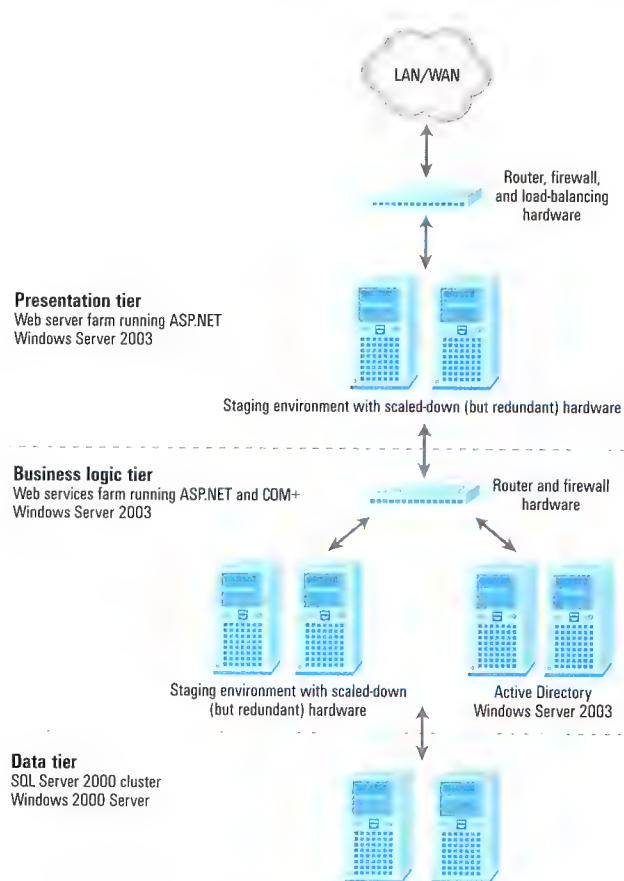


Figure 3. Staging environment for a three-tier architecture

in the document, no matter how trivial, must be added to it. The team should repeat the entire process until it is certain that the production rollout will be a success.

Deploying the .NET Framework

When rolling out a .NET application, the deployment team may encounter missing dependencies, such as the .NET Framework, code assemblies, COM objects, and third-party components. In addition, incorrect or inconsistent application configuration or security settings might affect deployment. Of these, the most significant and most common dependencies are the common language runtime (CLR) and the .NET Framework class library.

The CLR is the foundation of the .NET Framework. The CLR manages the code and objects at run time, providing core services such as memory management, thread management, and remoting services. Code that uses these core services through .NET is known as managed code. Components that do not target the CLR, but instead call directly to the operating system application programming interfaces (APIs), are known as unmanaged code. Managed code requires the .NET Framework on the physical machine on which it runs.

The .NET Framework class library is a collection of reusable classes provided as the basis for applications. These classes provide capabilities such as graphical user interface (GUI) elements, file I/O, data access, load balancing, message queuing, and transactions. Applications that make use of the .NET Framework classes require the .NET Framework on the physical machine.

The following is a list of machines in a typical development scenario and guidelines regarding whether they require the .NET Framework:

- **End-user machine:** End users generally access a Web application through a browser. Because the server machine runs the .NET-based code and passes down plain HTML to the browser, the .NET Framework typically is not needed on the client machine. An important exception is browser-based managed controls, which are assemblies referenced by Web pages that are downloaded to the user's computer and run locally. Browser-based managed controls are similar in concept to ActiveX® controls, except that they require the .NET Framework. When an application uses managed code, the .NET Framework must be installed on client computers.
- **Web servers:** Any .NET-based application that uses managed code must have the .NET Framework installed. Web server applications are no exception, and thus physical Web servers require the .NET Framework.
- **Business logic servers:** Many enterprise .NET applications implement a physically separate tier to provide business logic. If these applications use the .NET Framework to

provide or consume Web services, to access remote data stores, or to perform any other activity, these servers will need the .NET Framework installed.

- **Database servers:** Because database products, such as Microsoft SQL Server 2000, are native applications, they do not require the .NET Framework. This can be confusing, especially since they are often described as .NET Enterprise Servers.

The redistributable version of .NET (dotnetfx.exe) is required to deploy the .NET Framework to a server. This package is available from Microsoft and freely distributable to those with a licensed copy of Visual Studio .NET or the .NET Framework Software Development Kit (SDK). The SDK contains both the CLR and the .NET Framework classes.

Deploying the .NET Framework to older machines is not recommended. Improved performance and lower maintenance costs will likely offset the purchase price of buying new machines. Minimum requirements for existing hardware include an Intel® Pentium® processor at 133 MHz and 256 MB of RAM to run the .NET Framework on a server. For a client machine, minimum requirements are an Intel Pentium processor at 90 MHz and 96 MB of RAM. However, it is highly recommended that servers have at least an Intel Pentium 4 processor and 1 GB of RAM, and that client machines have at least 256 MB of RAM. Machines running Microsoft Windows NT® Server 4.0 must be upgraded to Microsoft Windows® 2000 Server or Windows Server 2003 to run ASP.NET.

Achieving successful deployments

Having an appropriate set of staging and production environments, in addition to a well-planned deployment process, is vital to a successful roll-out of .NET-based applications. Development and deployment teams should account for dependencies, configure staging servers properly, and document the deployment process meticulously. When upgrading production servers, the staging servers also should be upgraded so that the next deployment will run smoothly. Administrators who follow these techniques are more likely to achieve successful application deployments. ➤

Scott Jamison (scott_jamison@dell.com) is a technical strategist at Dell Professional Services. He is a Microsoft Certified Solution Developer for .NET and has been building solutions using Microsoft technology for more than 10 years. Scott has a master's degree in Computer Science from Boston University.

FOR MORE INFORMATION

<http://www.microsoft.com/.net>

Evaluating Windows Server 2003 Migration Options

Because Microsoft is discontinuing Microsoft® Windows NT® Server 4.0 and releasing Windows® Server 2003, organizations must soon decide whether to migrate the Windows-based Dell™ PowerEdge™ servers in their infrastructures to Windows 2000 Server or Windows Server 2003. Three migration options exist, depending on existing IT infrastructure, deployment time frame, desired operating system features, and organizational resources and goals.

BY MIKE OWENS

Most enterprises that use Intel® processor-based Dell™ servers in their IT infrastructures have systems running the Microsoft® Windows® 2000 Server operating system (OS) or Windows NT® Server OS. With the upcoming changes to Microsoft's product and support plans for the Windows NT Server OS, IT departments may soon need to consider an OS migration strategy for the Windows-based Dell PowerEdge™ servers in their infrastructures.

After five years on the market—and with the increasing adoption of its successor, the Windows 2000 Server family of operating systems—Microsoft is retiring Windows NT Server 4.0. Windows NT Server is currently no longer offered through direct original equipment manufacturer (OEM) channels, such as Dell, and by the beginning of 2005, Microsoft plans to have phased out support for the product.

In the spring of 2003, Microsoft is expected to launch Windows Server 2003, the follow-on OS to Windows 2000 Server. Built on the same architecture as Windows 2000 Server, the Windows Server 2003 OS offers new features that improve security, performance, administration, and ease of use.

Selecting an OS based on existing IT infrastructure

Both the retirement of Windows NT and the emergence of Windows Server 2003 will prompt organizations to evaluate their migration options as they search for the best paths for their IT infrastructures. The direction organizations decide to follow may largely depend on which OS they currently use.

Organizations currently running Windows NT Server

Organizations still running Windows NT Server are probably doing so for one of the following reasons:

- Certain legacy applications run only on Windows NT Server.
- Microsoft Active Directory® directory service migration is considered too large an undertaking.
- Organizational issues prevent funding or impede progress on migration projects.

Because of the support risks associated with remaining on Windows NT Server and the security improvements

in both Windows 2000 Server and Windows Server 2003, most organizations will elect to migrate despite these obstacles. Those that do face three primary migration options (see Figure 1).

Migrate to Windows 2000 Server. Following this migration path, organizations may choose to redesign their network and domain infrastructures and migrate to Active Directory. Alternatively, they may choose to simply have their Windows 2000 directory controllers emulate Windows NT directory controllers. Administrators will need to refresh server hardware and migrate key applications to systems running Windows 2000 Server. While some pockets of legacy applications may remain on Windows NT Server, the bulk of the infrastructure will run on Windows 2000 Server.

Maintain Windows NT domain and replace aging servers with Windows 2000 Server systems. By selecting this alternative, organizations maintain a legacy domain infrastructure based on Windows NT Server. They still retain the risk of losing support as Microsoft's phased retirement plan for Windows NT Server takes effect, and they are not able to take full advantage of the security enhancements available in Windows 2000 Server or Windows Server 2003 domains. However, this option is the least costly and complex. Because Windows NT Server is no longer on the market, administrators will replace aging servers with Windows 2000 Server systems through attrition and likely begin some application migration. Organizations that choose this path may migrate to either Windows 2000 Server or Windows Server 2003 at a later date. Dell recommends that organizations pursuing this option initiate migration plans as soon as possible.

Migrate to Windows Server 2003. Organizations that choose this option perform many of the same tasks that a full-scale migration to Windows 2000 Server would entail: redesign of the network and directory infrastructure, server refresh, and application migration. Windows Server 2003 offers the same basic architecture as

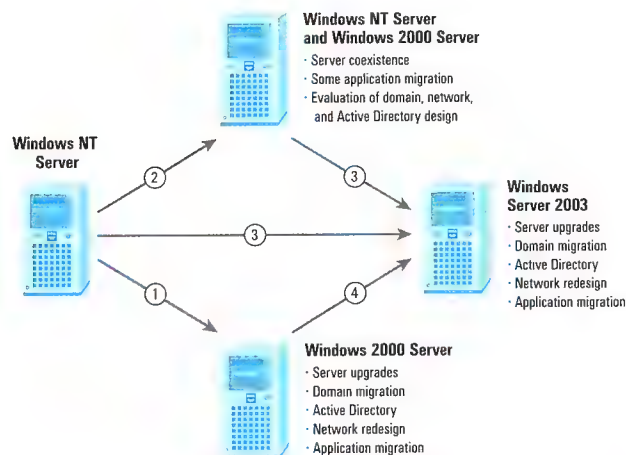


Figure 1. Migration options

IT departments may soon need to consider an OS migration strategy for the Windows-based Dell PowerEdge servers in their infrastructures.

Windows 2000 Server, but the newer OS provides a more robust set of features that eases some of the administrative burden associated with a large-scale migration.

Organizations currently running Windows 2000 Server

More than half of the Windows server installed base is already running the Windows 2000 Server OS. Windows 2000 Server on PowerEdge servers provides a

reliable and scalable platform that allows organizations to leverage the planning and development that they put into their Windows 2000 migration should they elect to migrate to Windows Server 2003 at a later date.

Organizations already running Windows 2000 Server can evaluate Windows Server 2003 to determine whether it better meets their IT and business needs. From a Windows 2000 Server environment, an upgrade to Windows Server 2003 is relatively straightforward—the most time-consuming work having been accomplished in the migration to Windows 2000 Server. Therefore, organizations can upgrade to Windows Server 2003 without great effort.

Considering migration time frame when selecting a new OS

Aside from the administrative burden of migration, organizations must also consider time frame when selecting a migration path. Groups currently evaluating, planning, or deploying a Windows 2000 Server infrastructure should continue with those efforts even if they someday plan to migrate to Windows Server 2003. The Windows 2000 Server and Windows Server 2003 operating systems are designed to coexist in the same environment, so the planning and effort invested in a migration to Windows 2000 remains relevant if an organization should later decide to adopt Windows Server 2003.

IT departments considering a migration to Windows Server 2003 so that they can develop and deploy Extensible Markup Language (XML) Web services using the .NET Framework and Microsoft Visual Studio® .NET development system can deploy the Windows 2000 Server family and be assured of a smooth migration to Windows Server 2003.

Organizations that have not begun planning a migration project should consider evaluating and deploying Windows Server 2003. Figure 2 provides a time line that can help serve as a guide for migrations from Windows NT Server.

Factoring OS features into the migration decision

As organizations consider migrating their servers from one Windows OS to another, they must factor their existing infrastructure

Dell server consolidation.

Saves money. Saves space.

Spells doom for your old servers.



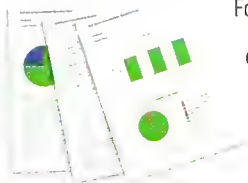


Consolidating servers with Dell's PowerEdge™ servers

Consolidate with Dell and you'll need to find a new use for your old servers.

What kind of server consolidation solutions does Dell bring to your enterprise? Just what you'd expect: A legendary focus on you, the customer, that's as relentless as our focus on driving down costs. An end-to-end solution that saves you money today and tomorrow by delivering:

- **Maximum flexibility, manageability, value and price/performance.** Our new line of PowerEdge™ servers, powered by Intel® Xeon™ processors, consistently ranks at the top of industry benchmarks such as TPC.* Collectively helping to lower TCO and accelerate time to a positive ROI.
- **Optimized uptime/maximized investment.** Dell's new systems management solutions deploy software, tools and services which help simplify and automate server systems administration. Leveraging your IT resources and maximizing your IT dollar.
- **Server infrastructure consolidation services.** Our comprehensive portfolio includes consolidation readiness assessment, consolidation design and transformation, customer training and certification, deployment and high availability support services.
- **Flexible financing alternatives.** Dell gives you a variety of financing avenues designed to help you optimize ROI.



For nearly 20 years, we've revolutionized the way the world buys and manages technology. Now find out how Dell's direct approach can revolutionize your server consolidation. To learn more about the Dell ROI test, visit www.dell.com/serverROI.



Flexible solutions that can cut costs today and tomorrow. Easy as **DELL**™

Visit www.dell.com/serverROI

*Per TPC-W 100,000 Item Config Results Test, July 2002. For more information, visit www.tpc.org. TPC-W is a trademark of the Transaction Processing Corporation. Dell, the Dell logo, and Xeon are trademarks or registered trademarks of Dell Corporation in the United States and other countries. Dell, the Dell logo, and PowerEdge are registered trademarks of the Dell Computer Corporation. ©2002 Dell Computer Corporation. All rights reserved.

and deployment time frame into the decision. OS feature sets provide an additional factor for consideration. The Windows Server 2003 OS offers new features that help organizations optimize and improve their Windows infrastructures. Figure 3 provides a high-level overview of the primary Windows Server 2003 editions that will be available and some of the basic feature specifications that differentiate these editions.

The Windows Server 2003 OS offers new features that improve security, performance, administration, and ease of use.

Evaluating business goals when considering a migration to Windows Server 2003

Built on the same basic architecture as Windows 2000 Server, the Windows Server 2003 OS running on PowerEdge servers offers the same scalability and reliability. However, a migration to Windows Server 2003 also imparts a longer product life cycle with the following additional benefits:

- **Active Directory improvements:** A primary improvement in Windows Server 2003 is the administration and management tools and features of Active Directory. The new features make migration, management, and administration of Active Directory easier and more efficient.
- **Server and storage management:** Windows Server 2003 includes features to ease the management of the server and storage resources in .NET infrastructures.
- **Security:** Windows Server 2003 was designed with a focus on security, and Microsoft has made several modifications

	Windows Server 2003, Web Edition (32-bit only)	Windows Server 2003, Standard Edition (32-bit only)	Windows Server 2003, Enterprise Edition (32-bit/64-bit)	Windows Server 2003, Datacenter Edition (32-bit/64-bit)
Processor support	Up to 2	Up to 4	Up to 8	Up to 64
Memory support	Up to 2 GB	Up to 4 GB	Up to 32 GB (32-bit) Up to 64 GB (64-bit)	Up to 64 GB (32-bit) Up to 512 GB (64-bit)
Failover high-availability (HA) clustering	N/A	N/A	8 nodes	8 nodes
Network Load Balancing	Up to 32 nodes	Up to 32 nodes	Up to 32 nodes	Up to 32 nodes
Terminal Services	Remote Desktop for Administration	Terminal Server and Remote Desktop for Administration	Terminal Server and Remote Desktop for Administration	Terminal Server and Remote Desktop for Administration
Windows System Resource Manager	N/A	N/A	Yes	Yes

Figure 3. Microsoft Windows Server 2003 OS features

not only to the tools and features of the OS but also to some of the default settings.

- **Reliability and performance:** Windows Server 2003 leverages the reliability of Windows 2000 Server, and also offers new clustering features as well as the ability to install larger cluster sizes.

With such a variety of reliable options available, organizations must consider the business drivers behind migration—such as IT resources and organizational goals—as well as current infrastructure, time factors, and desired OS features before solidifying a plan. ☞

Mike Owens (mike_owens@dell.com) is a Microsoft alliance manager with the Dell Product Group. Mike specializes in the Windows systems architecture. He has an M.S. in Information Systems from Drexel University and a B.S. in Computer Science from the United States Military Academy at West Point.

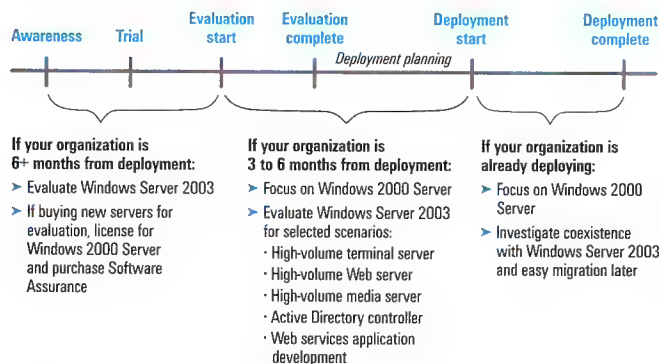


Figure 2. Migration time line

FOR MORE INFORMATION

Dell Windows Server 2003 migration services:
<http://www.dell.com/windows.net>

Windows Server 2003 preview version download or kit:
<http://www.microsoft.com/windowsserver2003>

New features in Windows Server 2003: <http://www.microsoft.com/windowsserver2003/techinfo/overview/reviewersguide.mspx>

Key retirement dates for Windows NT Server 4.0:
<http://www.microsoft.com/ntserver/ProductInfo/Availability/Retiring.asp>

Dell and Microsoft:

Partnering to Deliver Windows Server 2003 on PowerEdge Servers

To coincide with the final release of the Microsoft[®] Windows[®] Server 2003 family of operating systems, Dell is simultaneously releasing these operating systems on shipping Dell[™] servers and storage products. This article describes how Dell and Microsoft worked together to achieve a holistic release of Windows Server 2003 on Dell products.

BY LARA BENSON AND KEITH MATTESON

The Microsoft[®] Windows[®] Server 2003 operating system (OS) is a full-featured server OS that offers high levels of reliability, scalability, and security for enterprise computing environments. The final release of the Windows Server 2003 family will be available on all currently shipping Dell[™] PowerEdge[™] servers. Previous generation servers—including the third-generation PowerEdge 2300, 4300, 4350, 6300, and 6350 servers—will also support the new OS. In addition, the following Dell products will be compatible with Windows Server 2003:

- All currently shipping peripherals and most peripherals that were supported with previous generation servers
- All currently shipping RAID cards
- All currently shipping network interface cards (NICs), including available support for advanced features such as teaming

Although many peripherals are supported by native drivers in the Windows Server 2003 build, some newer peripherals may have updated drivers. Updates for these drivers are available through Dell Factory Install, through download from <http://support.dell.com>, and through Dell OpenManage[™] Server Assistant. The Dell OpenManage suite of management products will fully support Windows

Server 2003. Additionally, VERITAS Backup Exec[™] and Yosemite Technologies[™] TapeWare[®] tape backup and restore products will be available directly from Dell at or near the release date of the OS. The Dell release of Windows Server 2003 will also fully support storage area network (SAN) functionality for the Dell|EMC line of Fibre Channel storage products.

Achieving this holistic release of Windows Server 2003 on Dell server and storage products required a collaborative partnership from both Dell and Microsoft. The Dell and Microsoft program teams identified several areas in which continued contact and discussion would facilitate this release:

- **Technical issues:** Testing and issue management, beginning with the earliest original equipment manufacturer (OEM) release of Windows Server 2003
- **Program issues:** Candid discussions of non-technical issues that would inhibit the timely holistic release of Windows Server 2003 from Dell
- **Joint Development Program:** Support for Microsoft Joint Development Program (JDP) members deploying Windows Server 2003 on Dell hardware
- **Certification:** Participation in the creation of the “Designed for Windows” Logo Program specification and early engagement on known problem areas

By listening to customers, collaborating with Microsoft, and coordinating with each other, the Dell extended engineering and program management teams were able to deliver a holistic release.

Testing the technology

Testing, issues management, and issues resolution were critical to the success of the Windows Server 2003 comprehensive release. Both Dell and Microsoft wanted to ensure that thorough testing on all supported hardware and software from Dell occurred as early as possible. To accomplish that objective, all major releases of the OS underwent complete regression testing—testing only updates and changes—to validate resolution of the known issues and to develop a comprehensive test for any new issues.

Testing Dell hardware at Microsoft

Dell wanted to provide a broad test bed of the released Dell server hardware to the Microsoft server development organization, so that daily build testing would occur on production Dell hardware. Together with the Microsoft development teams, Dell identified components, specifications, and situations that would challenge performance of Windows Server 2003 on Dell servers. These areas included storage, networking, the Advanced Configuration and Power Interface (ACPI) specification, and stress testing. Dell then provided representative hardware, both to support the Windows Server 2003 project and to ensure that Dell hardware would be used in Microsoft labs. In certain cases, Dell also made prototype hardware available for early testing.

Daily builds of Windows Server 2003 were tested on Dell equipment, and issues that were identified by Microsoft, Dell, or Dell vendors during testing could be replicated quickly in the Microsoft labs on the existing equipment. This arrangement allowed all participants to more rapidly determine the root cause of issues and implement a resolution.

Implementing test plans at Dell

Dell participated in the Microsoft beta program and received major preproduction builds for testing. Dell tested these builds of the OS on production Dell equipment as well as on servers and software under development but not yet available to Microsoft. This testing of prototype servers with the prototype OS enabled early detection of problems that could be resolved for the final releases of each product. To test Windows Server 2003, the Dell operating systems engineering team assembled a test bed of all servers and peripherals that were to be supported with the new OS.

The testing teams. The operating systems engineering team comprised many subteams. Focus engineers were assigned for all major aspects of OS testing. The focus engineers wrote the test plans for their focus areas, executed the tests, worked with the extended engineering teams responsible for those products, and communicated with external test groups.

The interlock phase. Engineering teams within Dell adopted a formal, two-phase interlock process to prepare for the release of Windows Server 2003 on PowerEdge servers and storage. Dell defines an *interlock* as a series of meetings with a conclusion and a written agreement defining each team's deliverables, risks, and assumptions that apply to the release of the project. In this case, the interlock was between the Dell OS engineering team and each extended engineering team: RAID engineering, BIOS development, and server platform engineering.

The first interlock was based on the second beta release of the OS and involved the Dell OS engineering team and 46 internal Dell engineering teams. Completed early in the Windows Server 2003 release phase, this interlock process proved essential in identifying extended team schedules and risks.

For the second phase of interlocks, the OS engineering team met with 36 internal Dell engineering teams 10 months before the release of Windows Server 2003 from Dell. In addition to the two major interlock phases, re-interlocks were scheduled with all applicable engineering teams in response to each Microsoft schedule change. In this final phase, the teams compared the extended team schedules and risks to those identified in the initial interlock. This comparison showed the readiness of Dell to release Windows Server 2003.

Quick Test. As each build of the OS was received from Microsoft, the Dell team performed a Quick Test, a standard Dell functionality test, on a rack of representative servers that included a comprehensive mix of the supported peripherals. This Quick Test checked installation and basic functionality of all aspects of server operation and consumed between 24 and 48 hours. Performing the Quick Test helped to ensure that neither the extended teams nor the Dell OS engineering team would waste test cycles on builds that had major problems. Upon successful completion of the Quick Test, the builds were released to the extended engineering team, which consisted of all the development groups concerned with testing Windows Server 2003 on their products.

Regression testing. The team performed regression testing on interim builds between the major releases. This testing allowed Dell to track the progress on updates without waiting for the next major release to verify them.

Comprehensive testing. Full execution of the comprehensive test plan was scheduled for all three formal beta releases of Windows Server 2003. The test plan was updated and improved based on the issues found during the operating systems engineering testing. Issues identified by Microsoft, by the extended engineering teams at Dell, and by JDP customer feedback also helped improve the test plan. Adjustments to the test plan allowed for much more testing on the later builds—in the same amount of time.

By conducting tests at multiple sites using test teams that focused on different elements and followed different test plans, Dell reduced the risk of *test escapes*, or issues that should have been

Maximum ROI. Minimum IOU.



Dell | Enterprise

The power of Dell flexibility.

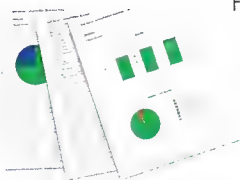
What does Dell bring to your enterprise? Just what you'd expect: A legendary focus on you, the customer, that's as relentless as our focus on driving down costs.

With Dell you get:

- **Enterprise solutions on your terms.** No proprietary systems, no endless consulting fees. Just powerful, cost-effective, industry-standard technology like Dell™ PowerEdge™ servers with Intel® Xeon™ processors.
- **Simplified systems management.** From SAN and server consolidation to UNIX® migration, we collate flexible systems that are easier to afford, year after year.
- **Enterprise services.** Dell has an expanded range of services to help you simplify the design and management of your IT infrastructure.



Dell PowerEdge servers use Intel Xeon processors



For nearly 20 years, we've revolutionized the way the world buys and manages technology. Now find out how Dell's direct approach can revolutionize your enterprise. To learn more about the Dell ROI test, visit www.dell.com/ROI1.



Enterprise solutions that can cut costs today and tomorrow. Easy as



Visit www.dell.com/ROI1

caught in the lab but were not. Furthermore, Dell could complete the comprehensive test passes and regression cycles more quickly and on many server models simultaneously.

Improved issue submission and tracking tools from Microsoft, and early and rigorous engagement with the Microsoft development team, contributed to effective issue tracking and resolution. Regular and comprehensive communication reduced last-minute surprises, and the “test early and test often” approach contributed greatly to the success of this release.

Coordinating the program

As well as communicating closely on technical issues, Dell maintained close program ties with Microsoft to help ensure a smooth product release. During quarterly on-site meetings, Dell and Microsoft identified areas of mutual interest and initiatives for cooperation. Action items from these meetings were tracked at regular conference calls. Dell and Microsoft also held periodic readiness meetings to discuss in detail all critical issues affecting the release date of Windows Server 2003. This approach allowed development teams in both companies to identify each other’s highest priority concerns and resolve the critical program and technical issues that came to light.

At Dell, the program management team helped drive the successful release of Windows Server 2003 through executive status meetings. In addition, weekly core team meetings involving representatives from marketing, engineering, procurement, finance, services, quality, and other groups helped identify program issues and risks early in the release.

Separating program concerns from technical testing and bug fixes helped Dell and Microsoft focus their resources and allowed for more day-to-day efficiency. Periodic discussions of all key program concerns helped both companies remove obstacles to the product release.

Engaging customers through the Joint Development Program

Through the partnership with Microsoft and the Microsoft JDP for early adopters of Windows Server 2003, Dell identified JDP members who were deploying on Dell equipment. Dell was able to offer support to these members and benefit from their feedback.

During this development phase, one key realization was that the third-generation PowerEdge servers (such as the PowerEdge 6300 and 6350) are still very much in use, and many JDP customer deployments of Windows Server 2003 are on third-generation hardware. Although Dell had not originally planned to support some legacy platforms for the holistic release, Dell revised its plans and now supports the third-generation servers. JDP feedback also led to extending Microsoft certification for third-generation servers, discussed in the next section.

Preparing for product certification

Certification, the final area of technical engagement, consists of compliance testing to determine whether servers, peripherals, and other components meet the requirements of the “Designed for Windows” Logo Program from Microsoft. Products that earn the Windows logo have been tested to ensure that they meet Microsoft standards for compatibility on the Windows OS designated on the logo.

The Dell certification team and the Windows Hardware Quality Lab (WHQL) team at Microsoft worked closely to discuss progress in certifying new server platforms and devices as well as to address issues identified during beta testing for certification under Windows Server 2003. The teams also collaborated to certify some legacy Dell servers under the Microsoft “Supported” program.¹

Achieving a holistic release

Several strategies helped Dell and Microsoft achieve a holistic release of Windows Server 2003 on Dell servers and storage products. To accomplish their goal, Dell and Microsoft cooperated to maintain early and regular communication at the technical and program management levels; to comprehensively test beta releases at Dell and incorporate Dell hardware into the test bed at Microsoft; to test deployment of prototype and production-level PowerEdge servers and hardware provided by Dell at the Microsoft Windows Server 2003 test labs; and to work with early adopters of Windows Server 2003 in Dell production environments. Articles in *Dell Power Solutions*, May 2003 Special Issue, and additional internal support tools released in parallel with Windows Server 2003, should further assist system administrators in maximizing the Windows Server 2003 OS on Dell hardware. ☞

Lara Benson (lara_benson@dell.com) is a project manager for Operating Systems and Cluster Engineering in the Dell Product Group. Lara has a B.A. in Business Administration from the University of Florida and a master’s certificate in Project Management from George Washington University.

Keith Matteson (keith_matteson@dell.com) is a senior software engineer in the Server Operating Systems Engineering Department in the Dell Product Group. He has a B.S. and an M.S. in Electrical Engineering from Cornell University.

FOR MORE INFORMATION

<http://www.dell.com/microsoft>

<http://www.microsoft.com/windowsserver2003>

<http://www.microsoft.com/windows>

¹ The “Supported” program is for previously certified systems that do not meet the full set of Windows logo program requirements (as defined in the Microsoft Windows Logo Program Desktop and Mobile PC Requirements, Version 1.1, and in the Windows Logo Program System and Device Requirements, Version 2.1), but which Dell continues to support because of strong customer demand.

Guide to Deploying Microsoft Windows Server 2003 on Dell PowerEdge Servers

Proper planning and installation information make deploying the Microsoft® Windows® Server 2003 operating system on Dell™ PowerEdge™ servers more efficient for system administrators. This article guides administrators through deployment planning; provides information about network, storage, and software components supported on PowerEdge servers running Windows Server 2003; and discusses procurement and installation options.

BY THE DELL SERVER OPERATING SYSTEMS ENGINEERING GROUP

The Microsoft® Windows® Server 2003 operating system (OS) brings new capabilities to Intel® processor-based servers, and Dell worked extensively with Microsoft throughout the development of the OS. During the development process Dell, its partners, and Microsoft performed comprehensive testing across Dell™ PowerEdge™ servers and other Dell hardware and software products to ensure compatibility with the new OS. In addition to testing performed in Dell labs, Dell worked with many customers in the Microsoft Joint Development Program (JDP). These customers received prerelease versions of device drivers, BIOS and firmware releases, and Dell OpenManage™ Server Administrator software components. The intent of the JDP was as follows:

- Help Dell and Microsoft customers to deploy the Windows Server 2003 family throughout the OS development process
- Capture customer issues from production environments

- Enhance Dell hardware and software product readiness for Windows Server 2003

Based on information gathered in the JDP and through extensive testing, this article discusses Microsoft Windows Server 2003 deployment planning, includes lists of supported hardware and software components, and explains installation options.

Deployment planning for the Windows Server 2003 OS

The information in this section is designed to help IT professionals deploy Windows Server 2003 on Dell PowerEdge servers. Deploying a new release of Windows Server 2003 requires careful planning and understanding of an organization's current IT infrastructure, supported systems, and devices and applications that will run under the new OS.

Proper planning begins with selecting the appropriate Windows Server 2003 edition for a server's intended

Edition	Physical memory supported	Number of processors supported	General features
Web	Up to 2 GB	Up to 2	Web hosting (Internet Information Services 6.0 and Microsoft ASP.NET), Internet Connection Firewall (ICF), Network Load Balancing (NLB), Distributed File System (DFS), Encrypting File System (EFS), Windows Management Instrumentation (WMI) command line, and .NET Framework
Standard	Up to 4 GB	Up to 4	All features of Web Edition plus domain controller, Microsoft Active Directory directory service, Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), Windows Internet Naming Service (WINS), Remote Installation Services (RIS), Terminal Server, and virtual private network (VPN) support
Enterprise	Up to 32 GB (32-bit version) Up to 64 GB (64-bit version)	Up to 8	All features of Standard Edition plus support for 8-node Cluster Service (MSCS), Terminal Server Session Directory, Microsoft Metadirectory Services (MMS), and Windows System Resource Manager (WSRM)
Datacenter	Up to 64 GB (32-bit version) Up to 512 GB (64-bit version)	Up to 32 (32-bit version) Up to 64 (64-bit version)	Supports Enterprise Edition software features and Datacenter services programs

Figure 1. Windows Server 2003 memory and processor support

Dell PowerEdge server	BIOS/ESM	Dell PowerEdge server	BIOS/ESM
PowerEdge 300	A03/NA	PowerEdge 2500	A06/A54
PowerEdge 350	A09/NA	PowerEdge 2550	A08/A56
PowerEdge 500SC	A07/NA	PowerEdge 2600	A04/A24
PowerEdge 600SC	A05/NA	PowerEdge 2650	A10/A02
PowerEdge 650	A01/NA	PowerEdge 4300	A12/A54
PowerEdge 1300	A12/NA	PowerEdge 4350	A06/A54
PowerEdge 1400	A09/NA	PowerEdge 4400	A10/A56
PowerEdge 1500SC	A05/NA	PowerEdge 4600	A08/A21
PowerEdge 1550	A08/NA	PowerEdge 6300	A13/A54
PowerEdge 1600SC	A04/NA	PowerEdge 6350	A13/A54
PowerEdge 1650	A09/A18	PowerEdge 6400	A13/A57
PowerEdge 1655MC	A00/NA	PowerEdge 6450	A13/A57
PowerEdge 1750	A00/A23	PowerEdge 6600	A09/A20
PowerEdge 2300	A15/A54	PowerEdge 6650	A09/A20
PowerEdge 2400	A09/A54	PowerEdge 7150	A04/A03
PowerEdge 2450	A08/A54	PowerEdge 8450	A06/A05

Figure 2. Supported PowerEdge servers and minimum supported BIOS and Embedded Server Management (ESM) versions

Product	Driver/location
Intel PRO/100 (82558)	E100b325.sys / Windows Server 2003 CD
Intel PRO/100+ (82559)	E100b325.sys / Windows Server 2003 CD
Intel PRO/100+ Dual Port (82558)	E100b325.sys / Windows Server 2003 CD
Intel PRO/100 S (82559)	E100b325.sys / Windows Server 2003 CD
Intel PRO/100 S (82550)	E100b325.sys / Windows Server 2003 CD
Intel PRO/1000 XF Intel PRO/1000 XT	E1000325.sys / Windows Server 2003 CD
Intel PRO/1000 F	E1000325.sys / Windows Server 2003 CD
Intel PRO/1000 T	E1000325.sys / Windows Server 2003 CD

Figure 3. Supported Intel PCI network adapters and their corresponding drivers

use (see Figure 1). The Windows Server 2003 family includes the Web Edition, Standard Edition, Enterprise Edition, and Datacenter Edition. The Microsoft Web site provides a comprehensive list of Windows Server 2003 features.

Assessing the existing IT environment

Collecting information about an organization's current IT infrastructure can be time intensive, but a well-documented infrastructure (including applications, BIOS and firmware revisions, and network components) is key to a successful OS migration. A clean, or "fresh," installation is always recommended. If an upgrade is planned, backing up critical data before installing the new OS is an important first step. Administrators should make a checklist of the applications installed on each server if they plan to set up the same infrastructure using the new OS. Finally, administrators must assess the servers and hardware components to determine which are supported on Windows Server 2003 and which must be replaced.

Supported Dell PowerEdge servers

Thorough testing by Dell and Microsoft engineering teams has helped to ensure the readiness of Windows Server 2003 on PowerEdge servers. Figure 2 provides a comprehensive list of the PowerEdge servers that support Windows Server 2003 and the lowest versions of BIOS supported. All the software and firmware versions specified in this document are the minimum versions supported by Windows Server 2003. Any later releases will also be supported by Windows Server 2003.

Dell-supported network components

The entire network architecture has a significant impact on the overall network performance of Windows Server 2003 systems. PowerEdge servers can include many peripherals, such as an embedded Peripheral Component Interconnect (PCI) network card that enables the system to interact with the physical network through the PCI interface. This section provides comprehensive lists of the Dell-supported peripherals that are compatible with Windows Server 2003. Figures 3 and 4 list supported Intel cards and compatible LAN on Motherboard (LOM) adapters, while Figures 5 and 6 list Broadcom® PCI network cards and compatible LOM adapters.

Product	Driver/location
PowerEdge 350, 500SC, 1400, 1500SC, 1550, 1600SC, 2400, 2450, 2500, 2550, 4400, 4600, 6400, 6450, 7150	E100b325.sys / Windows Server 2003 CD
PowerEdge 1650, 2600	E1000325.sys / Windows Server 2003 CD

Figure 4. Supported Intel LOM network adapters and their corresponding drivers

Product	Driver/location
BCM5700 10/100/1000BaseT	b57xp32.sys / Windows Server 2003 CD
BCM5701 10/100/1000BaseT	b57xp32.sys / Windows Server 2003 CD
BCM5703 10/100/1000BaseT	b57xp32.sys / Windows Server 2003 CD

Figure 5. Supported Broadcom PCI network adapters and their corresponding drivers

Product	Driver/location
PowerEdge 2550, 2650, 4600, 6600, 6650	b57xp32.sys / Windows Server 2003 CD

Figure 6. Supported Broadcom LOM network adapters and their corresponding drivers

Figure 7 lists the Windows Server 2003-compatible adapters that support teaming on Dell systems. Using teaming, administrators can improve throughput by balancing network traffic across a group of adapters. Administrators also can configure adapters so that if an adapter fails, a standby adapter can take over from the failed adapter and help prevent data loss.

Windows Server 2003 also provides software network bridging capability. A network bridge can create connections between different types of network media, help administrators to manage LAN segments, and create a single subnet for the entire network. No configuration or additional hardware is required.

Downloading device drivers

The latest BIOS, firmware, and device drivers may be obtained from the Dell support site (<http://support.dell.com>) or by using the latest Dell OpenManage Server Assistant CD that supports Windows Server 2003. The Dell OpenManage Server Assistant CD is a part of the Dell OpenManage kit that is shipped with every Dell PowerEdge server.

To obtain BIOS, firmware, or drivers from the Dell support site:

1. Visit <http://support.dell.com>.
2. Select the system to download by clicking the down arrow under "Select Server."

To obtain BIOS, firmware, or drivers using the Server Assistant CD:

1. Insert the Server Assistant CD into a machine running a Windows OS.
2. Select the system to update by clicking the down arrow under "Select Server."
3. Click the file to download under "Select Drivers/Utilities Set."
4. Click "Continue" to go to the next screen, and follow the

Product	Driver/location
Intel PRO family of adapters	ianswpx.sys http://support.dell.com or Server Assistant CD
Broadcom NetXtreme™ adapters	baspxp32.sys http://support.dell.com or Server Assistant CD

Figure 7. Supported advanced network teaming adapters and their corresponding drivers

instructions on the screen to download the BIOS, firmware, or driver.

Working with PCI devices using Device Manager

When the OS cannot load the necessary driver for a specific PCI device, administrators can update the driver using Device Manager, a user-level process that runs continuously and launches through the kernel:

1. Right-click the My Computer icon.
2. Select the Manage menu option. The Computer Management dialog box appears.
3. Select the Device Manager option.
4. Right-click the failed device.
5. Select the Properties menu option.
6. Select the Driver tab.
7. Click the Update Driver button. The Update Device Driver wizard appears. This wizard will help install the software for the failed device.
8. Follow the prompts until the proper driver is loaded.

Administrators can also obtain PCI device IDs through Device Manager.

Network devices not supported by Windows Server 2003

Some network devices will not be supported on the Windows Server 2003 operating system:

- 3Com® 3C980 network interface card (NIC), base drives only
- Basic Rate Interface (BRI) Integrated Services Digital Network (ISDN) cards
- Aztech modems
- Alteon® Gigabit Ethernet¹ NICs
- Giganet™ cLAN™ host bus adapters (HBAs) and switches

Dell-supported storage components

SCSI is the most common and predominant technology used to connect disk devices to servers. Fibre Channel technology has made significant inroads into the disk subsystem arena, especially for very

¹This term indicates compliance with IEEE® standard 802.3ab for Gigabit Ethernet, and does not connote actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

Product	Driver/location/firmware/BIOS
PERC 2	perc2.sys, version 5.2/ Windows Server 2003 CD <i>See important note under "Primary storage components."</i>
PERC 2/Single Channel (SC)	mraid35x.sys, version 5.2/ Windows Server 2003 CD
PERC 2/Dual Channel (DC)	mraid35x.sys, version 5.2/ Windows Server 2003 CD
PERC 2/Quad Channel (QC)	perc2.sys, version 5.2/ Windows Server 2003 CD
PERC 3/SC	mraid35x.sys, version 5.2/ Windows Server 2003 CD/1.72/3.27
PERC 3/DC	mraid35x.sys, version 5.2/ Windows Server 2003 CD/1.72/3.27
PERC 3/Dual Channel Lite (DCL)	mraid35x.sys, version 5.2/ Windows Server 2003 CD/1.72/3.27
PERC 3/QC	mraid35x.sys, version 5.2/ Windows Server 2003 CD/1.72/3.27

Figure 8. RAID controllers and their corresponding drivers

Product	Driver/location
PERC 2/Si PowerEdge 2400	perc2.sys, version 5.2/ Windows Server 2003 CD <i>See important note under "Primary storage components."</i>
PERC 3/Si PowerEdge 2450	perc2.sys, version 5.2/ Windows Server 2003 CD <i>See important note under "Primary storage components."</i>
PERC 3/Di PowerEdge 1650, 2500, 2550, 2650, 4400, and 4600	perc2.sys; perc2hib.sys, version 5.2/ Windows Server 2003 CD <i>See important note under "Primary storage components."</i>

Figure 9. Supported RAID on Motherboard (ROMB) controllers and their corresponding drivers

Product	Driver/firmware/location	Dell OpenManage driver/location
Cost Effective RAID Controller (CERC) ATA100/4CH	dellcerc.sys, version 6.25.32.0 (Firmware: 6.6.1) Windows Server 2003 CD	mraid35x.sys, version 6.30.2* Server Assistant CD

**To use Dell OpenManage Array Manager, administrators must update the driver dellcerc.sys, version 6.25.32.0 to mraid35x.sys, version 6.30.2 or later (which is the same driver family as the Dell PERC 2 and PERC 3 controllers).*

Figure 10. Supported Advanced Technology Attachment (ATA) controllers and their corresponding drivers



PCI Hot Plug is a standards-based technology that minimizes server downtime by allowing users to add or replace PCI devices while the operating system is running. Dell supports PCI Hot Plug functionality on the PowerEdge 4400, 4600, 6300, 6350, 6400, 6450, 6600, 6650, and 8450 servers, and additional hot-plug capable servers are in development. For more information about PCI Hot Plug and Windows Server 2003, see "Using PCI Hot Plug on Dell PowerEdge Servers Running Windows Server 2003" by Scott M. Callaway and Faisal Ahmed in *Dell Power Solutions*, May 2003 Special Issue.

Product	Driver/location
Ultra160 PCI controller	adpu160.sys, version 5.2/ Windows Server 2003 CD
Ultra 320 PCI controller	symmpi.sys, version 1.03.12 (Firmware: 1.00.03)/ Windows Server 2003 CD

Figure 11. Supported SCSI controllers and their corresponding drivers²

large environments and for storage area network (SAN) implementations. Dell offers primary, secondary, and external storage components that support Windows Server 2003.

Primary storage components

Figures 8 through 11 list storage components that are compatible with Windows Server 2003 and Dell systems along with the baseline versions of supported drivers.

IMPORTANT NOTE: In environments that have PowerEdge Expandable RAID Controller 2 (PERC 2), PERC 2/Single Channel Integrated (Si), PERC 3/Si, or PERC 3/Dual Channel Integrated (Di) with a controller firmware version earlier than 2.7-1, administrators must not use the native driver on the Windows Server 2003 CD. Instead, update the driver to at least version 2.7-1.

Although all the RAID and SCSI devices listed in Figures 8 through 11 work with the native drivers on the Windows Server 2003 CD (with the exception of PERC 2, PERC 2/Si, PERC 3/Si, and PERC 3/Di), Figure 12 lists recommended drivers and firmware versions that are essential when running the Dell OpenManage Array Manager systems management tool. (See important note above.)

Secondary storage components

Windows Server 2003 has a native backup utility to help protect data from accidental loss if a system experiences hardware or storage media failure. The backup utility can be used to create a duplicate copy of the data on the hard disk on another storage device. The backup storage medium can be a logical drive (such as a hard drive), a separate storage device (such as a removable disk), or an entire library of disks or tapes organized into a media pool and controlled by a robotic changer.

The backup utility creates a volume shadow copy of the data to create an accurate point-in-time copy of the contents of the hard drive, including any open files or files being used by the system. Users can continue to access the system while the backup utility is running without risking loss of data. Figure 13 lists the Dell PowerVault™ secondary storage components that function with the Windows Server 2003 backup utility.

Fibre Channel external storage components

Supported Windows Server 2003 external storage includes both Dell | EMC and PowerVault SANs.

²For all embedded and PCI-based Adaptec® SCSI controllers, use the native Windows Server 2003 drivers.

Product	Firmware	Driver
PERC 2/SC	3.13	6.32.2.32
PERC 2/DC on PowerEdge 2300, 4300, and 6300 servers	1.06	6.32.2.32
PERC 2/DC on PowerEdge 2400, 4400, and 6400 servers	1.06	6.32.2.32
PERC 3/SC, PERC 3/DC, PERC 3/DCL, and PERC 3/QC	1.78	6.32.2.32
PERC 4/SC and PERC 4/DC	3.15	6.32.2.32
PERC 4/Di on PowerEdge 1750 servers	4.0	6.32.2.32
PERC 4/Di on PowerEdge 2600 servers	2.23	6.32.2.32
PERC 4/im (embedded) on PowerEdge 1655MC servers	1.00.12.00	1.08.18
PERC 4/im 1.1	1.02.09.00	1.08.18
CERC ATA100/4CH	6.62	6.32.2.32
CERC ATA/2S	2.6.101011000	2.6.2002.1009
PowerEdge 2600 ROMB	2.01	6.25.32
PERC 2 PERC 3/Di (PowerEdge 2500 and 4600 ROMB) PERC 3/Di (PowerEdge 1650 ROMB) PERC 3/Di (PowerEdge 2650 ROMB)	2.7.1	2.7.1
PERC 2/Si (PowerEdge 2400 ROMB) PERC 3/Si (PowerEdge 2450 ROMB) PERC 3/Di (PowerEdge 2550 and 4400 ROMB)	2.7.1	2.7.1

*These firmware and drivers are posted to <http://support.dell.com> and are on the Dell Server Assistant CD 7.4

Figure 12. Minimum supported firmware and driver versions on Windows Server 2003

Dell | EMC SANs. For Dell | EMC Fibre Channel SANs and associated hardware and software components, support for Windows Server 2003 is offered with the Dell | EMC FC and CX series storage families. These families include the FC4500, FC4700, FC5300, CX200, CX400, and CX600 storage systems. More information on the supported configurations can be obtained from the EMC support matrices available at http://www.emc.com/horizontal/interoperability/interop_support_matrices.jsp.

PowerVault SANs. Figure 14 shows the Fibre Channel external storage components that will support Windows Server 2003 in the future. For additional information, administrators should contact their Dell representatives.

Storage devices not supported by Windows Server 2003

Some storage devices will not be supported on the Windows Server 2003 operating system:

- QLogic® 2100 Series HBAs
- PowerVault 530F storage array
- PowerVault 650F storage array
- PERC 1

Supported Dell software components

Windows Server 2003 is compatible with Dell OpenManage systems management products, which allow administrators to manage, monitor, and control the health of a Dell PowerEdge server from a

Product	Driver/location
PowerVault 100T DDS4 (digital data storage) tape drive	4mmdat.sys Windows Server 2003 CD
PowerVault 100T Travan40 tape drive	qic157.sys Windows Server 2003 CD
PowerVault 100T TR5 tape drive	miniqic.sys Windows Server 2003 CD
PowerVault 110T DLT1® (Digital Linear Tape) tape drive	dlttape.sys Windows Server 2003 CD
PowerVault 110T Linear Tape-Open® (LTO®)	ltotape.sys Windows Server 2003 CD
PowerVault 110T SDLT (Super Digital Linear Tape) tape drive	dlttape.sys Windows Server 2003 CD
PowerVault 110T SDLT320 tape drive	qntmsdlt.sys http://support.dell.com
PowerVault 120T DDS4 autoloader (Sony TSL-11000 autoloader)	ddsmc.sys Windows Server 2003 CD
PowerVault 120T DLT1 autoloader	adicsc.sys Windows Server 2003 CD
PowerVault 122T LTO VS80 tape library	powerfil.sys Windows Server 2003 CD
PowerVault 128T LTO, SDLT tape library	hpmc.sys Windows Server 2003 CD
PowerVault 130T tape library	Dell130.sys (Windows 2000 driver) http://support.dell.com
PowerVault 132T LTO, SDLT320 tape library	pv132t.sys http://support.dell.com
PowerVault 136T LTO, SDLT tape library	adicsc.sys Windows Server 2003 CD

Figure 13. Supported secondary storage components and their corresponding drivers

central or remote location. The Dell OpenManage suite can be downloaded from <http://support.dell.com> and is available through the Dell OpenManage Subscription Service. For more information about Dell OpenManage and Windows Server 2003, see “Deploying Dell OpenManage on Windows Server 2003” by Paul Rad; Mohammad Dhedhi; Joseph Ondrejick, Ph.D.; and Candace Hobson in *Dell Power Solutions*, May 2003 Special Issue.

Provisioning Windows Server 2003

To suit different customers' needs, Dell offers several methods for ordering and installing Microsoft operating systems on PowerEdge servers. Also, Windows Server 2003 supports many deployment

Product	Firmware/driver/location
Dell PowerVault 660F disk array	Firmware: 7.8.2-00 or later Dell PowerVault Fibre Channel Utilities CD, version 5.3
Dell PowerVault 51F, 56F Fibre Channel switches	Firmware: 2.6.0d or later Dell PowerVault Fibre Channel Utilities CD, version 5.3
Dell PowerVault 35F FC-to-SCSI bridge	Firmware: D9913D or later Dell PowerVault Fibre Channel Utilities CD, version 5.3
QLogic 2200 Series HBAs	Firmware: Ql2200.sys, version 8.00.09.24 or later Dell PowerVault Fibre Channel Utilities CD, version 5.3

Figure 14. Fibre Channel external storage to be supported on Windows Server 2003; firmware and driver versions are subject to change

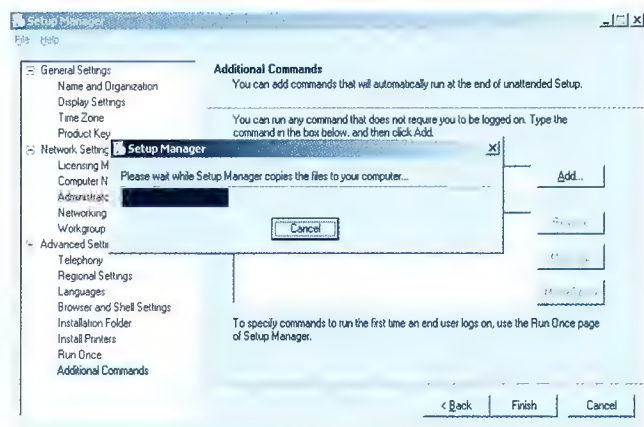


Figure 15. Setup Manager user interface

and support tools, which are available on the Windows Server 2003 CD in the support\tools directory. This directory also contains deptool.chm, an online help file. For more information about provisioning Windows operating systems, see "Provisioning Microsoft Operating Systems on Dell PowerEdge Servers" by Lara Benson and Matthew Paul in *Dell Power Solutions*, November 2002.

Dell factory installation (build-to-order process). When a customer orders factory installation of Windows Server 2003 on a PowerEdge server, Dell factory personnel incorporate the latest Dell-qualified and Dell-supported device drivers.

Installation through Dell OpenManage Server Assistant. Dell OpenManage Server Assistant helps administrators install and configure a supported version of Windows on a Dell PowerEdge server. The Server Assistant CD also contains drivers, system BIOS updates, and utilities that administrators need to operate PowerEdge servers. Server Assistant ships with all PowerEdge servers and also can be ordered through the Dell OpenManage Subscription Service.

Manual installation. Before manually installing the OS, administrators should download the latest version of the appropriate drivers from the Dell support site at <http://support.dell.com>.

Microsoft unattended installation. To save time when deploying multiple instances of Windows Server 2003, administrators can use Microsoft unattended installation,

Deploying a new release
of Windows Server 2003
requires careful planning
and understanding
of an organization's
current IT infrastructure,
supported systems, and
devices and applications
that will run under
the new OS.

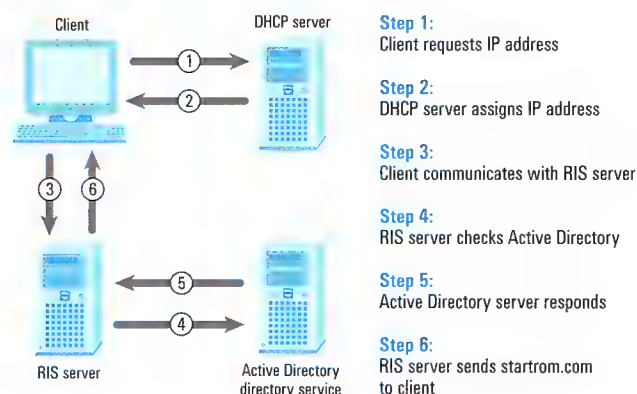


Figure 16. Using Windows Server 2003 Remote Installation Services (RIS) for server setup

which offers several tools that allow the use of answer files to automate the installation process. Using answer files, administrators can quickly install the Microsoft OS in unattended setup mode on multiple servers. Unattended setup does not require user intervention because the answer files contain all required setup information—including computer name and network adapter configuration—that administrators would normally enter during a standard installation.

Setup Manager (see Figure 15), which is located on the support\tools directory on the Windows Server 2003 CD, is a wizard-based deployment tool that assists administrators in automating either a clean installation or upgrade.

Remote Installation Services. Windows Server 2003 Remote Installation Services (RIS), which permit network-initiated setup, are capable of deploying both client and server operating systems onto a bare metal server (see Figure 16). One of the significant improvements in Windows Server 2003 is that RIS allows administrators to deploy server operating systems, including Windows Server 2003 (except Datacenter Edition) and Windows 2000 Server (except Datacenter Edition and 64-bit versions).

Target machines must meet the minimum hardware requirements for the Windows Server 2003 OS, support the Preboot Execution Environment (PXE) option in their system BIOS, and have at least one supported network adapter that complies with PXE Dynamic Host Configuration Protocol (DHCP)-based boot ROM version 1.0 or later. In addition, the boot floppy disk supports certain additional PCI network adapters that are not PXE capable.

A list of supported network adapters can be found in the rbfex.exe utility available in the Windows Server 2003 installation CD.

Microsoft System Preparation tool. The Microsoft System Preparation (Sysprep) tool allows administrators to take a snapshot of one system and copy it to an identical system (see Figure 17). This is a particularly fast way of deploying a fully configured system—including software and OS settings—to multiple target systems. Certain restrictions apply when using Sysprep. In

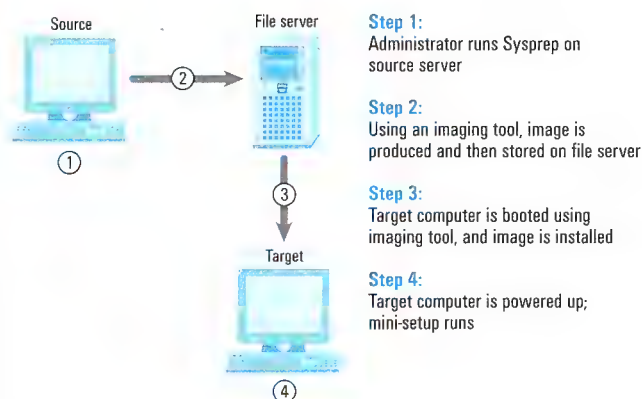


Figure 17. How Sysprep works

particular, source and target systems should be identical in the following aspects:

- Advanced Configuration and Power Interface (ACPI) support
- Hardware abstraction layers (HALs)
- Mass-storage device controllers

Plug-and-play devices need not be the same because Sysprep detects these devices on the first reboot.

Windows Server 2003 installation options

While a clean installation is always recommended, this article also provides information about upgrading. Administrators should be sure to back up critical data before installing the new OS when performing an upgrade.

Performing a Windows Server 2003 clean installation

To successfully perform a clean installation of Windows Server 2003 on a PowerEdge server (see Figure 18), perform the following steps sequentially:

1. **Prepare BIOS and firmware.** A clean installation starts with an upgrade to the Dell server's BIOS, Embedded Server Management (ESM), and primary storage systems. Before installing Windows Server 2003, administrators should download the appropriate BIOS and ESM from the Dell



Figure 18. Overview of the clean installation process

support site (<http://support.dell.com>) or use the latest Server Assistant CD that supports Windows Server 2003.

2. **Install the Windows Server 2003 OS.** Administrators can perform a clean installation by using the Server Assistant CD. Administrators also can perform a clean installation manually by using the retail CD or System Locked Preinstallation (SLP)-enabled CD.
3. **Verify device drivers.** After installing the Windows Server 2003 OS, administrators must use Device Manager to verify that all device drivers have installed with no problems and are working correctly. An improperly working device is usually an indication that a driver is needed. In such an instance, administrators can download drivers from the Dell support site or the Server Assistant CD.
4. **Install Windows Server 2003 versions of Dell software.** Administrators finish the upgrade process by installing Windows Server 2003 versions of Dell tools and software, such as Dell OpenManage components.

Upgrading to Windows Server 2003

To successfully upgrade to Windows Server 2003 on a PowerEdge server (see Figure 19), perform the following steps sequentially:

1. **Uninstall previous versions of Dell software.** An upgrade installation starts with an uninstallation of previously installed Dell tools and software, such as Dell OpenManage components.
2. **Upgrade BIOS and firmware.** Administrators should upgrade the server's BIOS, ESM, and primary storage firmware. Before installing Windows Server 2003, administrators should download the appropriate BIOS and firmware from the Dell support site (<http://support.dell.com>) or use the latest Server Assistant CD that supports Windows Server 2003.
3. **Upgrade to the Windows Server 2003 OS.** An upgrade to the Windows Server 2003 OS is supported only through the retail CD.
4. **Install Windows Server 2003 versions of Dell software.** Administrators finish the upgrade process by installing Windows Server 2003 versions of Dell tools and software, such as Dell OpenManage components.



Figure 19. Overview of the upgrade installation process

Windows Server 2003 activation: WPA versus SLP

The Windows Server 2003 OS must be activated after installation. An OS installed manually using a Microsoft retail CD is activated through Windows Product Activation (WPA), which requires each installation of the OS to be activated either online or by phone through a Microsoft License Server clearinghouse.

The Windows Server 2003 CD that ships with PowerEdge servers has a built-in anti-piracy technology known as System Locked Preinstallation (SLP). The SLP feature enables administrators to bind the OS to a system's specific hardware so that activating Windows Server 2003 is not necessary. When an SLP-enabled CD is used to install the OS, administrators need not type in a unique product key.

Because SLP-enabled CDs are designed only for clean installations of Windows Server 2003, administrators installing the OS using the CD should also boot from it. SLP is not supported while running setup.exe or winnt32.exe, because these executable files run from within an existing Windows environment.

An SLP implementation is transparent to the end user, without any noticeable difference from a manual installation using retail media. However, the SLP process works only on supported PowerEdge servers that ship with Windows Server 2003. In addition, any tampering with the SLP-enabled CD automatically invokes

WPA. The SLP-enabled CD is available only for 32-bit versions of Windows Server 2003, not 64-bit versions.

Smooth deployment and product compatibility

The compatibility testing of Dell servers as well as network, storage, and software components has helped Microsoft and Dell to ensure a holistic release of Windows Server 2003 on PowerEdge servers. The article "Dell and Microsoft: Partnering to Deliver Windows Server 2003 on PowerEdge Servers" by Lara Benson and Keith Matteson in *Dell Power Solutions*, May 2003 Special Issue, contains detailed information about the collaboration between Dell and Microsoft during the development of the new OS. The compatibility that resulted from this close partnership, along with the recommended installation and migration techniques discussed in this article, can facilitate a fast and efficient deployment. ☺

FOR MORE INFORMATION

<http://www.microsoft.com>

<http://support.dell.com>



An Opportunity to
Spread Your Message...

Dell Power Solutions delivers a portfolio of knowledge, practical experience, and best practices from Dell.

Written by Dell™ engineers and guest experts, the articles emphasize how current technology issues relate to today's enterprise and data center computing needs.

Those interested in advertising in *Dell Power Solutions* should visit <http://www.dell.com/powersolutions> and click on the Advertise icon for more information and a current media kit.

Deploying Dell OpenManage

on Microsoft Windows Server 2003

Dell™ OpenManage™ systems management products help administrators manage and configure Dell PowerEdge™ servers. The new Microsoft® Windows® Server 2003 server operating system offers ease of deployment, a secure infrastructure, high performance, and greater productivity. This article describes deployment and migration strategies for Dell OpenManage version 3.4 or above on Windows Server 2003.

BY PAUL RAD; MOHAMMAD DHEDHI; JOSEPH ONDREJIK, PH.D.; AND CANDACE HOBSON

Hardware failures on a server running mission-critical applications are costly and time-consuming. Proactive monitoring of system health can help prevent disasters. Dell™ OpenManage™ systems management software includes Server Administrator, a tool that monitors the health of individual servers and detects potentially problematic components, helping to prevent hardware failures. Dell OpenManage systems management products are critical to the management and configuration of Dell PowerEdge™ servers.

During development of the Microsoft® Windows® Server 2003 operating system (OS), Dell performed comprehensive testing of its OpenManage products on all new and previously shipped PowerEdge servers to help ensure full compatibility of Dell OpenManage products with Windows Server 2003. Dell also provided prerelease versions of Dell OpenManage to selected early Windows Server 2003 customers through the Microsoft Joint Development Program (JDP). This program helped Dell customers deploy or migrate to Windows Server 2003, and customer feedback helped Dell to enhance its hardware and software product features and readiness for Windows Server 2003.

Dell OpenManage offers administrators many systems management capabilities for Dell hardware running Windows Server 2003. Dell simplifies the installation of

systems management on systems running this new OS, and several migration strategies are available to ease the transition from a previous Windows OS to Windows Server 2003.

Running Dell OpenManage systems management products on Windows Server 2003

Dell OpenManage allows administrators to manage, monitor, and control the health of Dell PowerEdge servers from a central or a remote location. The software includes a management console to monitor networked servers and integrated functions to manage them.

Management console for many servers

A centralized network management console allows administrators to manage one or more servers remotely. Dell OpenManage IT Assistant performs system discovery, asset management, asset reporting, event monitoring, and remote configuration of Dell PowerEdge servers, including those running Windows Server 2003 or previous releases of Windows operating systems.

In Figure 1, the management console runs IT Assistant to centrally monitor and manage systems on the network. When an error is discovered on a managed server, IT Assistant receives notification of the event. IT Assistant can be configured to send e-mail or page the system

OBTAINING DELL OPENMANAGE

All Dell PowerEdge servers ship with three CDs that contain systems management tools and documentation:

- **Dell OpenManage Server Assistant CD:** Contains drivers and utilities that Dell PowerEdge servers need in order to operate, and helps install a supported version of Windows Server 2003 on a PowerEdge server.
- **Dell OpenManage Systems Management CD:** Includes Server Administrator, IT Assistant, and other associated Dell systems management software and utilities for Dell PowerEdge servers.
- **Dell Product Documentation CD:** Provides manuals and user guides for Dell PowerEdge servers, systems management software products, and peripherals.

The Dell OpenManage suite also can be downloaded from <http://support.dell.com>, and is available from the Dell OpenManage Subscription Service. An annual subscription includes a minimum of four updates to Dell systems management software.

administrator about this event, helping to prevent disasters. These capabilities extend management flexibility and increase availability of an IT environment. IT Assistant also allows a centralized launch point for other Dell OpenManage applications, such as Server Administrator, Array Manager, and Remote Access.

Managed servers for increased uptime

Servers typically can be managed directly or from a central console. Dell OpenManage Server Administrator facilitates the management of individual Dell servers from any location and provides access to the in-depth instrumentation built into a PowerEdge server. When Server Administrator detects a potential problem, it sends an event to management consoles such as IT Assistant for further analysis by the system administrator. Dell OpenManage Server Administrator provides the following services on Windows Server 2003:

Instrumentation service. Server Administrator offers a low-level hardware instrumentation service that supports a variety of network management protocols, including Simple Network Management Protocol (SNMP), Common Information Model (CIM), and Windows Management Instrumentation (WMI). This service provides access to hardware, BIOS, and firmware for proactive monitoring and alerting of server health to prevent disasters; presents configuration options for custom thresholds and BIOS options directly through Server Administrator or SNMP; and supports systems management consoles from major vendors (such as Tivoli, Hewlett-Packard, and Computer Associates) through the Dell OpenManage Connections program.

Windows Server 2003

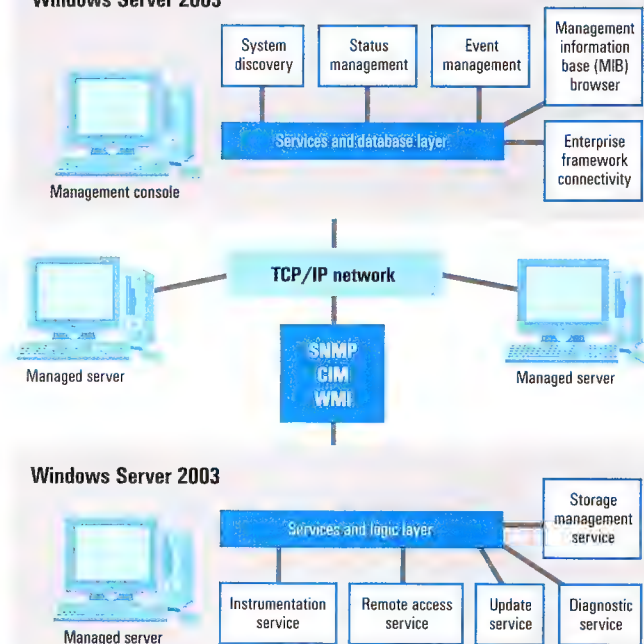


Figure 1. Running Dell OpenManage systems management products on Windows Server 2003

Remote access service. The remote access controller service minimizes costly downtime by supporting remote maintenance and control facilities for failure analysis and recovery. This service provides console redirection of and constant access to a system, even when that system is powered down.

Update service. The update service provides up-to-date version control and a valuable change management tool for performing BIOS and firmware updates on the local system.

Diagnostic service. The diagnostic service helps diagnose problems that affect the various hardware components of a PowerEdge server, including the CPU, memory, SCSI controller, and Peripheral Component Interconnect (PCI) bus.

Storage management service. This service uses Dell OpenManage Array Manager software to offer an integrated storage management solution. Through a common interface, Array Manager helps administrators configure and manage local and remote storage attached to a server while the server is online and continuing to process requests. When Array Manager is installed as part of Server Administrator, the Array Manager data is displayed within Server Administrator.

Dell has enhanced the functionality of Array Manager by incorporating the capabilities that have been added to the logical disk manager (LDM) on Windows Server 2003.

Dell OpenManage Remote Install. Remote Install is a software tool that lets administrators remotely deploy OS software to Dell PowerEdge 1655MC servers. Remote Install performs unattended, simultaneous provisioning and configuration of systems over the

network by using image-based technology. The imaging process involves capturing an image from a source system, creating an image file in a storage location, and then remotely deploying the image to multiple target systems. Images are managed by an image server, which is the system running Remote Install. Administrators can manage the imaging process remotely from a central console—a system running a Web browser and displaying content served by the image server.

Performing a clean installation of Server Administrator on Windows Server 2003

Figure 2 provides an overview of the clean installation process.

Upgrade the system BIOS and firmware. Verifying and upgrading the system BIOS and firmware to the recommended version for an IT organization's system and OS is an important step before installing Windows Server 2003. The latest versions of BIOS, firmware, and drivers for individual subcomponents that support the Windows Server 2003 OS on Dell PowerEdge servers are available at <http://support.dell.com> or from the latest Dell Server Assistant CD, which supports Windows Server 2003. To download BIOS and firmware from the Server Assistant CD, administrators can follow these steps:

1. Insert the Server Assistant CD into a system running the Microsoft Windows OS.
2. Select the system to be updated by clicking the down arrow under "Select Server."
3. Select "System Utilities" under "Select Drivers/Utilities Set."
4. Click "Continue" to go to the next screen and click the BIOS, firmware, or both (if either or both require an update) and save the file to the local hard drive.
5. To create diskettes to update the BIOS, firmware, or both, run the downloaded executable file.

Install Windows Server 2003. The Windows Server 2003 release of Dell OpenManage Server Assistant installs a supported version of the Windows Server 2003 OS on a PowerEdge server. The Server Assistant CD contains drivers and utilities that PowerEdge servers need so they can operate. Server Assistant also simplifies the installation of Windows Server 2003 on RAID-capable storage as well as other peripheral devices.

Verify the system using Device Manager. After installing Windows Server 2003 on a PowerEdge server, administrators should use Device Manager to verify that all device drivers are installed and working properly. If the device is not working properly, Device Manager indicates that a driver is needed.

Install the Windows Server 2003 version of Dell OpenManage. After completing system verification, administrators can set up Dell OpenManage using the Systems Management CD that is designed for Windows Server 2003. These tools facilitate installation of Server Administrator and IT Assistant.

Migrating Dell OpenManage products to Windows Server 2003

To migrate Dell OpenManage products to Windows Server 2003, administrators should begin by assessing the existing environment. Although time-consuming, collecting information about the current management infrastructure is necessary to plan a successful systems management migration. A checklist of previously installed Dell OpenManage contents on each server is recommended for administrators planning to maintain the same systems management infrastructure. Before migration, administrators should collect the following information:

- IT Assistant database
- Remote access service login and password
- Instrumentation login and password
- CIM and SNMP protocol attributes
- Custom or non-default instrumentation service threshold settings
- Scheduled diagnostics
- Server Administrator preferences
- IP addresses, if Dynamic Host Configuration Protocol (DHCP) is not used

Supplied with the checklist information, administrators can begin the first step in migrating Dell OpenManage to Windows Server 2003 (see Figure 3). Backing up crucial data on each server helps facilitate a system recovery if problems occur during the upgrade process. A complete system backup allows administrators to retrieve data stored on each system.

In addition, Dell strongly recommends that all managed server products be uninstalled before upgrading to Windows Server 2003. The architecture of the Dell OpenManage products revised for Windows Server 2003 has been optimized for that OS, and previous

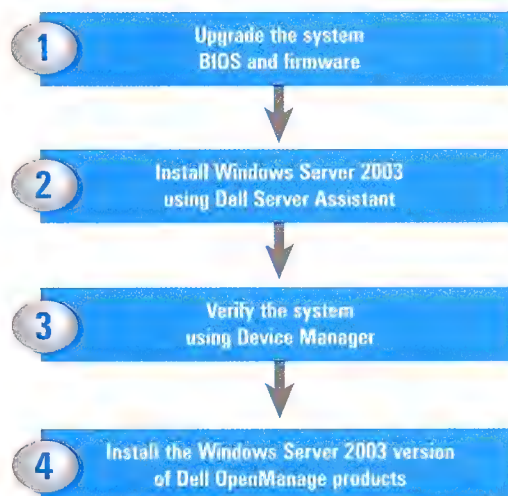


Figure 2. Clean installation of Dell OpenManage Server Administrator on Windows Server 2003



Figure 3. Migration of Dell OpenManage products to Windows Server 2003

versions of Dell OpenManage are not fully compatible with the Windows Server 2003 OS architecture. Figure 4 provides a comprehensive list of previously released Dell OpenManage products. Administrators using older Dell Server Agent software must uninstall Server Agent last, because application dependencies exist between other Dell OpenManage products and Server Agent. During the upgrade process, Windows Server 2003 might display an upgrade report that shows whether a previously installed program or driver will work after the upgrade. Dell also recommends uninstalling these applications or drivers before upgrading.

Maximizing Windows Server 2003 on PowerEdge servers

Obtaining maximum benefit from new, reliable, and enhanced operating systems such as Windows Server 2003 requires reliable servers as well as enhanced management tools. The management tools can minimize downtime and proactively monitor the health of the infrastructure to help avoid unexpected failures.

While Microsoft was developing the Windows Server 2003 OS, Dell developed OpenManage for Windows Server 2003 to incorporate new features and enhancements. Dell performed comprehensive testing on Dell OpenManage for both new and previously shipped PowerEdge servers to help ensure full compatibility with Windows Server 2003.

Previously released Dell OpenManage products	Dell OpenManage version 3.4 supported on Windows Server 2003
Array Manager	Yes
Server Administrator	Yes
Server Agent	Yes (consolidated within Server Administrator)
Version Assistant	No
Dell Remote Assistant Card (DRAC II)	Yes
Dell Remote Access Card (DRAC III)	Yes
Resolution Assistant	No
PowerEdge Expandable RAID Controller 2 (PERC 2)	Replaced by Array Manager
LSI Logic™-supplied PERC console package	Replaced by Array Manager
Local update	Yes
Online server diagnostics	Yes
IT Assistant	Yes
Remote Install (supported on the PowerEdge 1655MC)	Yes (supported on the PowerEdge 1655MC)

Figure 4. Support under Dell OpenManage version 3.4 for previously released Dell OpenManage components

Deploying a new release of an OS requires careful planning, knowledge of the current infrastructure, and a clear understanding of the supported systems, devices, and applications under the new OS. Dell is committed to helping administrators migrate to Windows Server 2003 as quickly as possible. ☺

Paul Rad (paul_rad@dell.com) is a senior software engineer in the Server Operating Systems Engineering Department in the Dell Product Group. Paul has master's degrees in both Computer Science and Computer Engineering from the University of Texas at San Antonio.

Mohammad Dhedhi (mohammad_dhedhi@dell.com) is a software engineer in the OpenManage Software Installation Group at Dell. Mohammad has a B.S. and an M.S. in Computer Science from The University of Texas at Austin.

Joseph Ondrejik, Ph.D. (joseph_ondrejik@dell.com) is a senior software engineering manager in the Dell Product Group—Software Development and is responsible for the Dell OpenManage Server Administrator instrumentation service. He has a Ph.D. in Computer Systems Integration from Colorado University.

Candace Hobson (candace@rice.edu) is an intern in the Server Operating Systems Engineering Department in the Dell Product Group. She is working on her B.S. in Computer Engineering at Rice University.

FOR MORE INFORMATION

<http://www.dell.com/openmanage>

Join the Discussion

about next-generation solutions for today's enterprise

Power Solutions Webcasts and seminars are comprehensive peer-to-peer technology forums bringing people together to discuss the best practices and technologies that are addressed on the pages of *Power Solutions* magazine. IT managers, technology leaders, infrastructure architects, and key decision makers can join the discussion, either by watching a Webcast online or attending a seminar on location.

Presented by *Power Solutions* contributors, and industry experts, Webcasts and seminars take place throughout the year and cover various topics. The following schedule lists upcoming *Power Solutions* Webcasts and seminars:

WEBCAST SERIES

Microsoft® .NET—The Next-Generation OS

Time: Each Wednesday at 12 P.M. Central Standard Time (CST)

May 14	June 11
May 21	June 18
May 28	June 25
June 4	

SCHEDULED TOPICS

Windows® Server 2003
Windows NT 4.0 Migration/Server Consolidation
Microsoft SQL Server™ 2000
Microsoft Exchange Server 2003
.NET Application Deployment and Strategy
Migration to Dell™ PowerEdge™ Environment



SEMINARS

Building Better Business Agility with
Microsoft .NET Solutions

June 2003–December 2003

For more information about *Power Solutions*
Webcasts and seminars, visit
http://www.dell.com/powersolutions_webcasts
and
http://www.dell.com/powersolutions_seminars



Using Windows System Resource Manager

for Server Consolidation

The Microsoft® Windows® System Resource Manager (WSRM) resource management tool is available for use with the Windows Server 2003 operating system, Enterprise Edition and Datacenter Edition. This article examines WSRM as a server consolidation tool and describes how this tool can be used to manage the allocation of resources, such as CPU and memory, across applications on a Dell™ PowerEdge™ server.

BY RANJITH PURUSH, CHRIS STACKHOUSE, AND PAUL RAD

When several applications run on a single server, the applications compete for system resources. Poor resource allocation can worsen the situation by causing bottlenecks, resource starvation, and an eventual decline in overall system performance. To prevent resource contention between applications, IT administrators can use resource management software, partitioning, or both.¹

Resource management software can increase server utilization by giving administrators control over the system resources available to individual applications, services, and processes within the operating system (OS). Administrators can set targets for the amount of resources consumed by particular applications or services. This supervision helps prevent applications from interfering with each other. For example, running two CPU-bound

applications at the same time could cause one of them to starve for CPU resources. Resource management helps avoid such situations.

A resource management tool, the Microsoft® Windows® System Resource Manager (WSRM), is available in the Microsoft Windows Server 2003 OS, Enterprise Edition and Datacenter Edition. This article discusses resource management on a Dell™ PowerEdge™ 6650 server running Windows Server 2003, Enterprise Edition. The PowerEdge 6650 is well suited to the demanding requirements of business-critical computing. In this scenario, the server is running two common applications: Microsoft Internet Information Services (IIS) 6.0, and SQL Server 2000. A Dell PowerEdge server ordered with Windows Server 2003, Enterprise or Datacenter Edition, includes a WSRM CD as part of the operating system kit.

¹ For more information about resource management and partitioning, see "Server Consolidation: Examining the Application Consolidation Approach" by Paul Rad and Dwayne Rodi in *Dell Power Solutions*, November 2002.

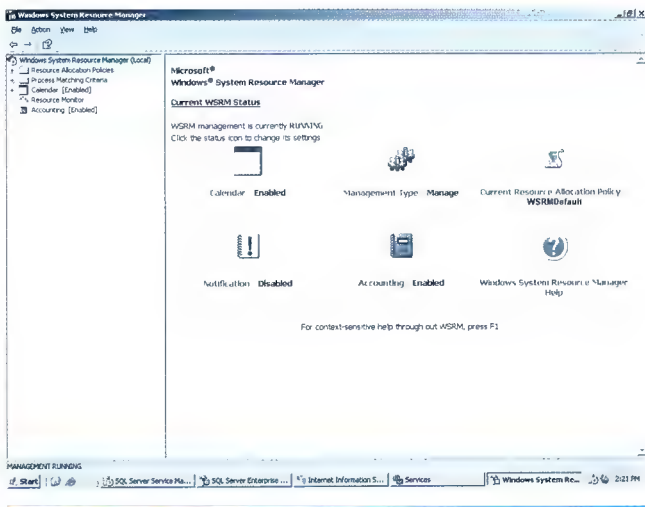


Figure 1. The Windows System Resource Manager GUI

Because this article was written before the release of Windows Server 2003, the WSRM tool used in the article is a prerelease version. For the latest information and updates on WSRM, visit: <http://microsoft.com/windowsserver2003/downloads/wsrn.msp>.

Introducing WSRM

WSRM allows administrators to control CPU and memory resource allocation to applications, services, and processes. This feature can be used to manage multiple applications on a single computer or multiple users on a computer that runs Microsoft Terminal Services. The WSRM architecture also allows administrators to manage resources on multiple systems. WSRM provides a graphical user interface (GUI), as shown in Figure 1, as well as a command-line interface for managing server resources.

Managing CPU resources

In a WSRM environment, processes are categorized into two types: managed and unmanaged. Managed processes are those whose resource utilization is controlled by WSRM. Unmanaged processes, typically system processes, are those that are not controlled by WSRM. Unmanaged processes have the highest priority during OS scheduling and can consume as much of the system resources as they need. WSRM can manage only CPU resources that are not being consumed by unmanaged processes. Therefore, the managed available CPU resources might be less than the total CPU resources available in the system.

Using WSRM, CPU resource management can be achieved either by allocating CPU resources or by setting processor affinity. In the former case, administrators can set an upper limit to the percentage of total CPU resources (all processors in the system) that can be utilized by a process. In the latter, administrators can assign processes to run on a specific processor or a specific range of processors.

The resource allocation limits are strictly enforced when multiple processes contend for CPU resources. However, in the absence of contention, managed processes can use more than their allocated CPU resources. Tests performed on the Intel® Xeon™ processor-based Dell PowerEdge 6650 server revealed that this flexibility in WSRM was especially advantageous when the workloads on the SQL Server and IIS Web server varied.

Although resource allocations between the two applications were made based on average workloads, at certain times heavier loads on the Web server existed while the database requests were less than average. WSRM permitted the IIS Web server to utilize more than its allocation resources by picking up the unused resources allocated to the SQL Server service.

Managing memory

WSRM allows administrators to control two aspects of memory utilization by a process: the process's working set size and its committed memory size. A process's working set memory is a subset of its virtual pages that are resident in physical memory. When the upper limit is defined for a process's working set size, WSRM prevents that process's working set size from exceeding the limit.

Setting such limits can be especially advantageous when the system has limited memory or when applications are very memory intensive. When the process's working set size reaches the limit, WSRM initiates paging of its working set, thus preventing any subsequent application errors.

WSRM also allows an upper limit to be set on the amount of committed memory that a process consumes (see Figure 2). Committed memory refers to the physical memory for which space has been reserved on the disk paging file in case it must be written back to disk. This type of memory either resides in the working set of a process or maps to a disk file. When a process increasingly consumes committed memory, a memory leak is usually to blame. Whenever a process consumes more than the allocated amount of committed memory, WSRM can either terminate the application or log an error in the system event log (SEL).

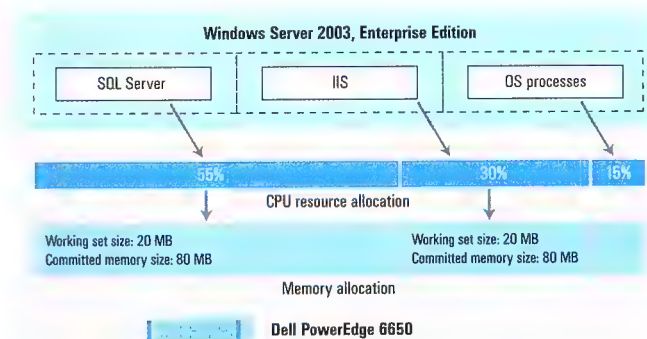


Figure 2. Example of WSRM using resource-allocation policy to manage resource utilization

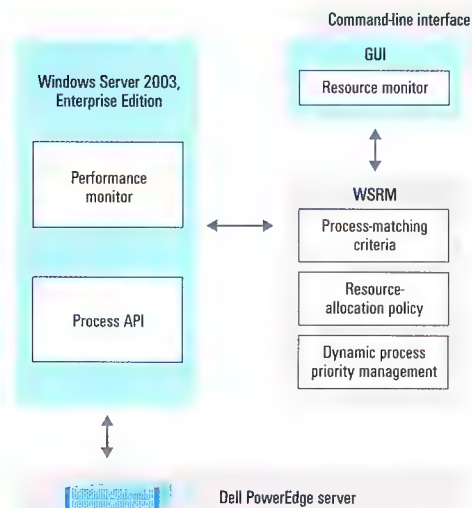


Figure 3. WSRM architecture

Understanding the WSRM architecture

WSRM manages and controls system resource usage by using a dynamic process-priority management algorithm that determines how to allocate resources among multiple processes. The WSRM service monitors the resource allocations at a regular polling interval. During each poll, the service determines whether the processes match their target allocations. If a process exceeds its target allocation, the WSRM service lowers the priority of the process.

A process consists of one or more threads, and the thread is the executable component of the process. Because the priority of the parent process greatly influences the priority of its threads, lowering the priority of a process also lowers the priority of its threads in most cases. The Windows OS implements a priority-driven,

CALENDAR AND ACCOUNTING FUNCTIONS

Two particularly useful features of WSRM are the following:

- **WSRM calendar:** This feature allows automatic activation of different resource-allocation policies based on a predefined schedule. The calendar is especially useful when the load on managed applications, and hence their resource requirements, vary systematically over a period of time.
- **WSRM accounting:** This function provides a mechanism to store and retrieve records of the behavior of managed processes. This information can be used to generate reports, provide charge-back accounting data, and also to study application behavior from the standpoint of resource utilization.

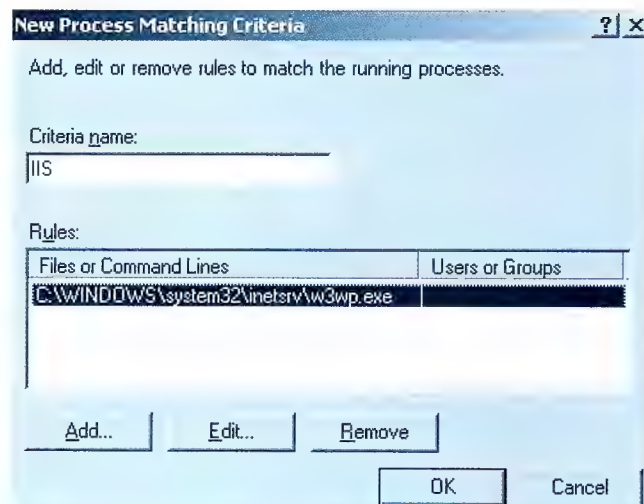


Figure 4. Creating a process-matching criterion

preemptive (round-robin) scheduling system that executes threads based on their priority—with higher priority threads executing first. Therefore, lowering the priority of a process indirectly reduces its resource utilization.

Users can control the resource utilization of applications, services, and processes through the two main components of WSRM: process-matching criteria and resource-allocation policies (see Figure 3).

Using process-matching criteria to define managed processes

WSRM uses a process-matching mechanism to group multiple processes that match administrator-specified rules. This group is then assigned to a resource-allocation policy. This policy sets the system resource limits for that group.

Figure 4 shows a process-matching criterion that matches the processes handling IIS 6.0 requests. The IIS 6.0 architecture introduced the concept of application pools that allow Web site applications to be isolated from each other. This feature can be leveraged by WSRM to set resource limits for applications in each pool. This action helps a single server to host multiple Web sites while ensuring that no single application will starve any others of resources.

Some critical system processes, such as the local security authority service, are by default set to be unmanaged, so that these processes can utilize as much of the system resources as required. This default setting is implemented in WSRM by including these processes in the system-defined exclusion list. Administrators can also set other processes to be unmanaged. Although the system-defined exclusion list itself cannot be modified, administrators can add or remove processes to or from the user-defined exclusion list. All processes in the system-defined exclusion list and in the user-defined exclusion list become unmanaged processes.

Using resource-allocation policies to specify usage levels

The resource-allocation policy is the component of WSRM that specifies resource-allocations for managed processes. A resource-allocation policy consists of one or more process-matching criteria and one or more of the following for each criterion: an associated CPU target, a memory limit, or a processor affinity. Any modifications to either the process-matching criterion or the resource-allocation policy will take effect immediately without any special action such as system reboot.

In the example shown in Figure 5, the policy named IIS_SQLServer allocated resources to two process-matching criteria. The IIS criterion matched the processes that were handling IIS 6.0 requests and the SQLServer criterion represented the SQL Server service. In this case, the resource-allocation policy reserved 35 percent of CPU resources for IIS and 55 percent for SQL Server. This allocation was designed to meet the Web and database workloads that were very specific to the client-server configuration that was set up for the scenario in this article.

After resources have been allocated to a group of processes, the WSRM service monitors the resource usage of each managed process in the group. If a process exceeds its resource allocation, the WSRM service attempts to bring usage back to its specified target. To accomplish this adjustment, WSRM changes the process priority to affect the OS schedule that determines which applications have access to CPU resources.

When a resource-allocation policy controls a group of processes that match a given process-matching criterion, the allocated processor resources for that criterion are divided among the matched processes. However, the memory resources are allocated on a per-process basis to each matched process within the group (see Figure 6).

Using WSRM, CPU resource management can be achieved either by allocating CPU resources or by setting processor affinity.

by the default group. The default group has resources not yet allocated to managed processes and unused by unmanaged processes.

Using the Resource Monitor to view performance data

The Resource Monitor is the GUI for WSRM performance monitoring. This GUI allows administrators to collect and view real-time performance data from a local computer or several remote computers. The Resource Monitor is a standard Windows performance monitor

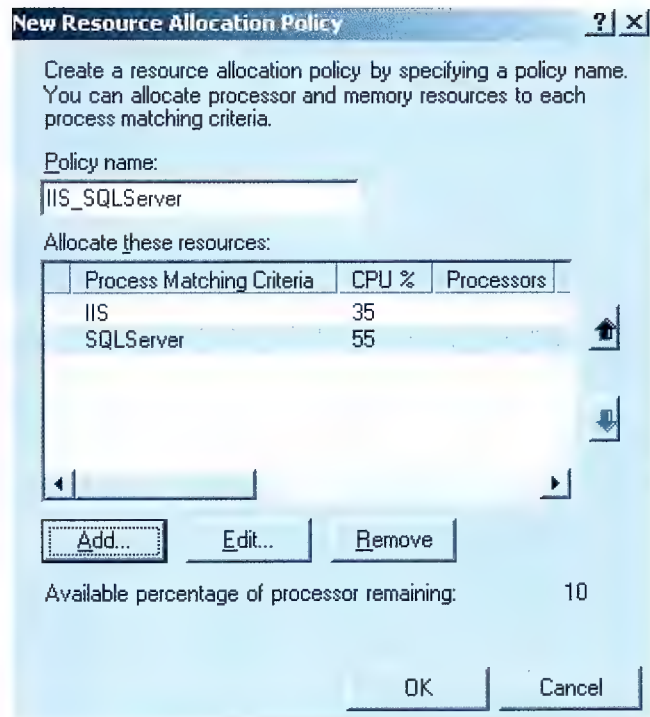


Figure 5. Creating a resource-allocation policy

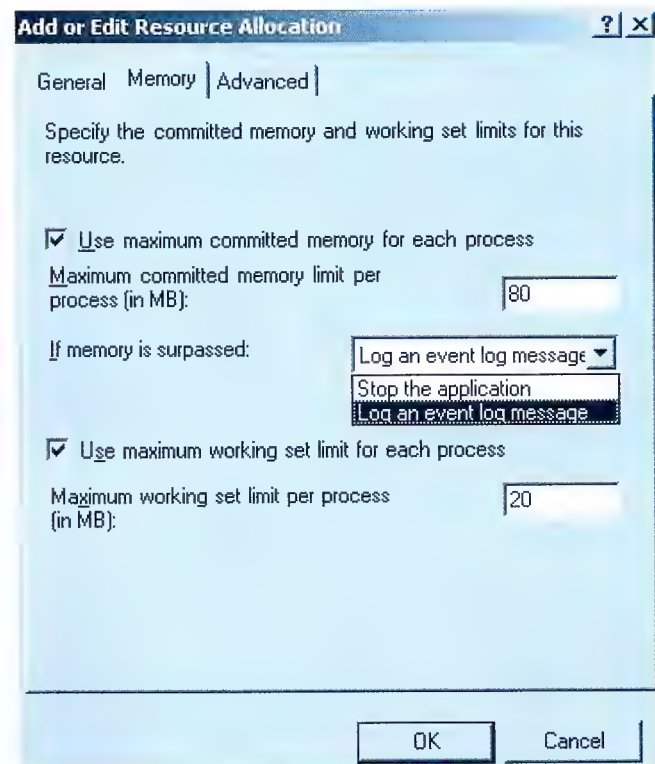


Figure 6. Example of memory allocation for a specific group of processes

Using WSRM can help prevent resource contention, bottlenecks, and other detriments to system performance.

(PerfMon) that has three additional performance counters—WSRM: Policy, WSRM: Process, and WSRM: Process Matching Criteria.

Figure 7 shows the resource utilization by the SQL Server and IIS processes for the scenario examined in this article. Although SQL Server was allocated only 55 percent of CPU resources, it used more than its allocation. The upper limit restriction for resource utilization is implemented by WSRM only when resource contention occurs. In this scenario, the only other running process, IIS 6.0, already had access to its allocated 35 percent of CPU resources, thus avoiding any contention.

When all the processes that match a process-matching criterion are stopped, the WSRM service reassigns the criterion's CPU allocation to the other process-matching criterion for management. Figure 8 shows the results when the IIS Web server was shut down. The process-matching criterion for SQL Server then managed nearly 100 percent of system CPU resources. This result allowed the SQL Server process to utilize most of the processing power in the system if the need arose. However, Figure 8 shows that the SQL Server load required only about 80 percent of available CPU.

Increasing server utilization

These tests on an Intel Xeon processor-based Dell PowerEdge server showed that when several applications are run on a single server, WSRM can effectively increase server utilization. WSRM

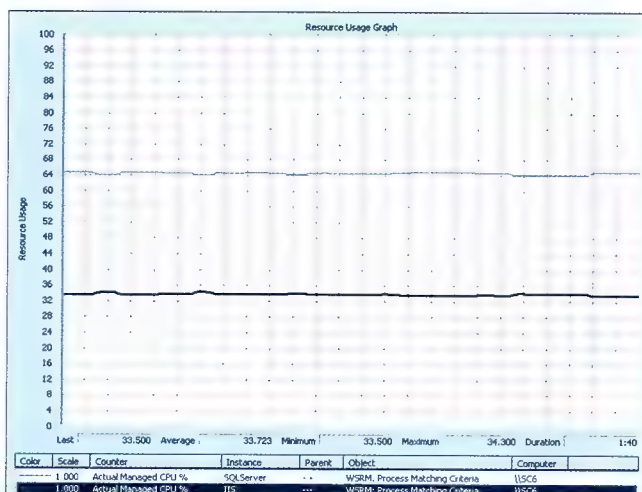


Figure 7. WSRM Resource Monitor showing CPU utilization when CPU allocation was 55 percent for the SQL Server process and 35 percent for IIS

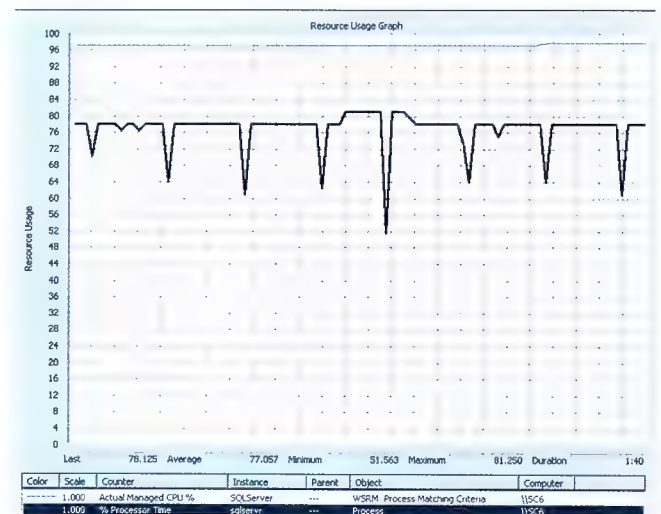


Figure 8. CPU utilization by the SQL Server process when IIS was shut down

can achieve these gains by giving administrators control over the system resources available to individual applications, services, and processes within the OS. Using WSRM can help prevent resource contention, bottlenecks, and other detriments to system performance. ➔

Ranjith Purush (ranjith_purush@dell.com) is a systems engineer in the Server Operating Systems Engineering Department at Dell. Ranjith has an M.S. in Electrical and Computer Engineering from The University of Texas at Austin.

Chris Stackhouse (cstack@microsoft.com) is the program manager for Microsoft Windows System Resource Manager. Chris has a degree in Electrical Engineering from the University of Washington and a master's degree in Computer Engineering from the University of Arizona.

Paul Rad (paul_rad@dell.com) is a senior software engineer in the Server Operating Systems Engineering Department in the Dell Product Group. Paul has master's degrees in both Computer Science and Computer Engineering from the University of Texas at San Antonio.

FOR MORE INFORMATION

WSRM:

<http://www.microsoft.com/windowsserver2003/downloads/wsrml.msp>

Examining Windows Server 2003

Performance on the Dell PowerEdge 4600 Server



The Microsoft® Windows® Server 2003 operating system provides networking and Web server enhancements, along with new features that improve overall system performance. To determine the benefits of upgrading to Windows Server 2003 from previous Microsoft operating systems, this article explores the performance of Intel® Xeon™ processor-based Dell™ PowerEdge™ 4600 servers running two industry-standard benchmarks that measure Web server performance and file and networking performance.

BY QUYNH O. NGUYEN AND MIN-JOHN LEE

Interest in upgrading to the Microsoft® Windows® Server 2003 operating system has grown because Windows Server 2003 takes advantage of the latest hardware advances, such as Intel® Hyper-Threading technology, non-uniform memory access (NUMA), and double data rate (DDR) memory, as well as the bandwidth boost in chipsets. Hardware technology can significantly affect the performance of server applications. For example, Hyper-Threading processors have the ability to execute multiple threads simultaneously on each physical processor.¹

To compare the difference in Web and networking performance across three evolutionary versions of the Windows operating system, this article reports the results of benchmark tests on identically configured Intel Xeon™ processor-based Dell™ PowerEdge™ 4600 servers running Microsoft Windows Server 2003, Enterprise Edition; Microsoft Windows 2000 Advanced Server; and Microsoft Windows NT® Server 4.0, Enterprise Edition. The test programs focused on the Web server, file services, and networking services that are core functions for an operating system kernel—in contrast to add-on services such as

Microsoft Exchange Server or Microsoft SQL Server, which have their own data engines. The tests were performed by Dell labs in May 2002.

In general, Windows Server 2003, Enterprise Edition improves upon the performance of Windows 2000 Advanced Server and Windows NT Server 4.0, Enterprise Edition because it has full awareness of the hardware architecture. For example, Windows Server 2003 distinguishes between the physical and logical processors of Hyper-Threading systems (see Figure 1), which gives server applications such as Microsoft Exchange Server and SQL databases a boost.

Configuring the systems for benchmark testing

Two benchmarks were used in the Dell study: Ziff Davis® WebBench® 4.1, which measures Web server performance, and Ziff Davis NetBench® 7.0.2, which measures file and networking performance. To run WebBench and NetBench, three Dell PowerEdge 4600 servers were configured with the same hardware and software, as listed in Figure 2. The configurations were based on the three generations of

¹ For more information on Hyper-Threading technology, see "An Introduction to HyperThreading Technology in the Intel Xeon Processor Architecture" by Humayun Khalid, Ph.D., in *Dell Power Solutions*, August 2002.



Windows under test. Although the Dell PowerEdge 4600 has a maximum memory capacity of 12 GB, Windows NT Server 4.0, Enterprise Edition can accommodate only 4 GB of memory. To maintain consistency across all the test platforms, each system was configured with only 4 GB of memory. Also, the /3GB switch was not turned on in the boot.ini file because this switch does not have the same effect in the three different operating systems.

Windows NT Server 4.0, Enterprise Edition limits the paging file to 4095 MB, so the Dell team configured the servers according to this limitation. To gain system efficiency, the paging file was moved to a separate partition, sitting on RAID-5 hardware configured with each server's on-board PowerEdge Expandable RAID Controller 3, Dual Channel Integrated (PERC 3/Di).

The test systems ran the Microsoft implementations of the Web server: Internet Information Server (IIS) 4.0 on Windows NT Server 4.0, Enterprise Edition²; Internet Information Services (also IIS) 5.0 on Windows 2000 Advanced Server; and Internet Information Services 6.0 on Windows Server 2003, Enterprise Edition, which is built into the operating system.

Setting up the network and testing infrastructure

The Ziff Davis benchmark suite uses as many as 60 client systems that simulate a real-world scenario to measure throughput and overall server response times. During the tests, multiple clients stressed each server with a large number of concurrent requests. WebBench tests throughput and requests per second, which are the most important parameters for Web servers. This benchmark also can simulate and measure performance for environments using static content and dynamic content. NetBench reports throughput and response time, measuring how well a file and network server processes file I/O requests from many Windows clients over the network.

	Windows NT Server 4.0, Enterprise Edition	Windows 2000 Advanced Server	Windows Server 2003, Enterprise Edition
Recognizes Hyper-Threading logical processor	No	Yes	Yes
Recognizes Hyper-Threading logical processor as physical processor	No	Yes*	No*
Distinguishes between Hyper-Threading logical processor and physical processor	No	No	Yes
Optimizes for Hyper-Threading	No	No	Yes

*Windows 2000 Advanced Server does not distinguish between logical and physical processors; it counts both types of processor against the license limit. Windows Server 2003 counts only the physical processors against the license limit.

Figure 1. How Microsoft Windows operating systems recognize Hyper-Threading processors

Server system	Dell PowerEdge 4600
CPU	Dual Intel Xeon processors at 2.2 GHz
System BIOS	A03, with Hyper-Threading enabled
System firmware	Embedded Server Management (ESM) firmware: 1.35; primary backplane firmware: 0.25
System memory	4 GB of 200 MHz DDR SDRAM
RAID controller	PERC 3/Di, BIOS: 2.7
Operating system	Installed in a mirrored 4 GB RAID-1 volume with two Seagate® ST336706LC hard drives
Data store	101.7 GB allocated in data store volume with four Fujitsu® MAM3367MC hard drives
OS paging file	4095 MB allocated in data store volume
On-board Ethernet adapter	One Broadcom® NetXtreme™ Gigabit Ethernet controller

Figure 2. Hardware configuration for the three Dell PowerEdge 4600 servers under test

To meet NetBench and WebBench requirements, the Dell team configured the servers, controllers, and clients into three distinctive sets:

- **Servers:** The target server was the PowerEdge 4600, running one of the three Microsoft operating systems: Windows NT Server 4.0, Enterprise Edition; Windows 2000 Advanced Server; or Windows Server 2003, Enterprise Edition. The server executed requests sent from the clients.
- **Controllers:** A controller running Windows 2000 was positioned as the mediator between the target server and its clients. The controller's primary job was to manage loads and generate requests for the clients (the clients in turn sent the requests to the target server for actual processing). A test suite specified how the controller should run each test, and each test suite was connected to several detailed script files that were run during the server stress tests. In addition, the controller recorded all test results and generated reports dynamically according to each test suite.
- **Clients:** As many as 58 clients, running different versions of compatible Windows operating systems, generated requests to the server under test. These requests to the target server were triggered by the controller's execution of a particular test suite specified in the benchmark scripts. A shared folder in each operating system was mapped to every client machine running NetBench.

Clients and servers were networked with Gigabit Ethernet³ connections and were located in different subnets. Three hubs

² The default for Windows NT Server 4.0, Enterprise Edition is IIS 3.0, but security and performance improvements between versions 3.0 and 4.0 are substantial enough to warrant upgrading from IIS 3.0 to IIS 4.0 on Windows NT Server 4.0, Enterprise Edition.

³ This term indicates compliance with IEEE® standard 802.3ab for Gigabit Ethernet, and does not connote actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

existed between clients and the target server because the test lab used a two-tier network design, in which one Extreme Networks® Summit7i™ switch served as a top-tier, central switch connecting to the second-tier, local switch in each subnet.

When traveling from a client to a server, data packets first arrived at the Dell PowerConnect™ 3024 local switch, which utilized two Gigabit Ethernet ports for this test. Next, the data packets were routed to the central Summit7i switch. In turn, the central switch distributed the data packets to the local switch that was connected to the target server.

Analyzing Web server performance

The two key indicators for Web server performance were number of requests per second and throughput in bytes per second. WebBench stressed the server until its performance started to saturate, meaning that the server could not efficiently process any more traffic, regardless of how many additional requests were sent. The analysis considered network capacity, to make sure that the performance saturation was not caused by a network bottleneck.

Test results indicate that the Windows NT Server 4.0 server had the slowest performance among the three operating systems tested. For that reason, the Windows NT Server 4.0 server became the baseline from which the performance differences were calculated for the other two operating systems (see Figure 3). The data shows that Windows 2000 Advanced Server and Windows Server 2003, Enterprise Edition do have a significant overall performance gain over Windows NT Server 4.0, Enterprise Edition. However, WebBench results consistently show no appreciable performance gain between Windows 2000 Advanced Server and Windows Server 2003.

To determine the differences between Windows 2000 Advanced Server and Windows Server 2003, Enterprise Edition, the Dell team

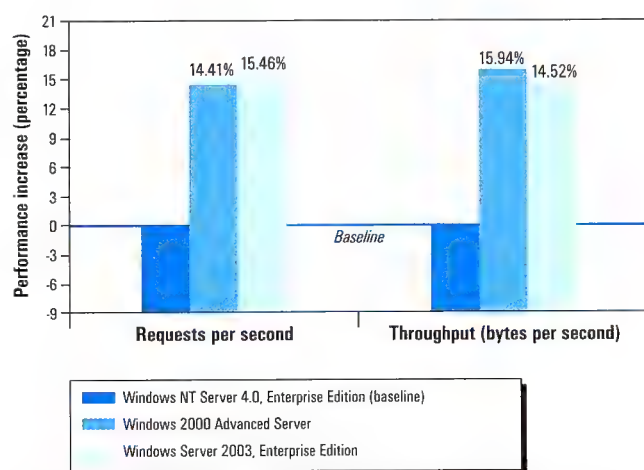


Figure 3. WebBench performance of Windows Server 2003, Enterprise Edition and Windows 2000 Advanced Server, using Windows NT Server 4.0, Enterprise Edition as the baseline

Counter name	Counter description for Windows Server*
\Memory\Pages/sec	The number of pages read from or written to disk to resolve hard page faults
Processor[_Total]\%Processor Time	The percentage of time that the processor is executing a non-idle thread
PhysicalDisk[_Total]\%Disk Time	The percentage of elapsed time that the selected disk drive is busy servicing read or write requests

*Source: Windows Server help file

Figure 4. Hardware utilization performance counters for Windows 2000 Advanced Server and Windows Server 2003, Enterprise Edition

examined the performance monitor counters that were recorded during the server test runs. Figure 4 shows the performance counters that recorded differences in hardware utilization.

Figure 5 shows the significant differences in hardware utilization among the three operating systems, using Windows NT Server 4.0, Enterprise Edition as the baseline performance comparison. To execute the same amount of tasks, the Windows Server 2003, Enterprise Edition server required only 54.50 percent of the CPU cycles and 3.41 percent of the disk I/O. In comparison, Windows 2000 Advanced Server required more computation cycles, more memory, and more disk activities: 70.99 percent of the CPU usage and 10.49 percent of the disk time. These results indicate that Windows Server 2003 is more suitable for consolidating different server applications into one physical server—or fewer physical servers—than Windows 2000 Advanced Server. Windows Server 2003 provides the better option for enterprises considering server scalability.

Assessing network performance

The two key indicators for networking performance were throughput in megabits per second and average response time in milliseconds. During the NetBench simulation, each client recorded the number of bytes moved and the amount of time needed to move those bytes. The number of bytes was then divided by the amount of time to calculate throughput—the total megabits per second for all clients. Figure 6 shows the throughput results for each operating system.

To measure the average response time, each client kept track of the number of times it executed an I/O call and the amount of time

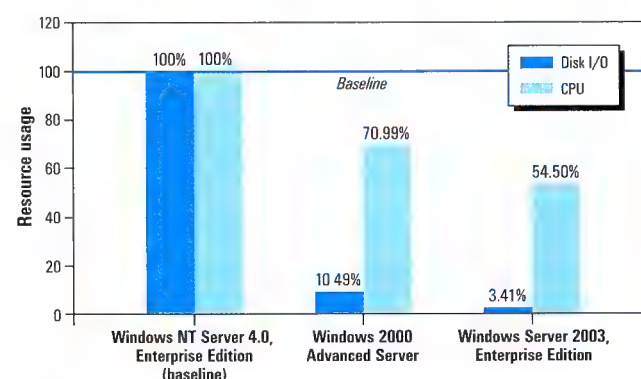


Figure 5. Web server CPU and disk I/O utilization during WebBench tests

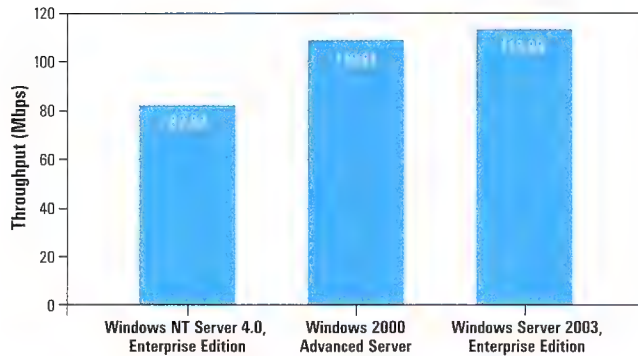


Figure 6. NetBench throughput results

each call took. The average frequency and duration of all the calls was then calculated to determine the average response time. Figure 7 shows that Windows Server 2003, Enterprise Edition clocked an average response time of 1.70 milliseconds, 71.4 percent faster than the response time for Windows NT Server 4.0, Enterprise Edition, which provided an average response time of 5.94 milliseconds. In turn, Windows Server 2003 was 46.4 percent faster than Windows 2000 Advanced Server, which had an average response time of 3.17 milliseconds. The impact of the average response time performance gain for Windows Server 2003 over its predecessors is even more significant considering its superior throughput. Given the same average response time, Windows Server 2003, Enterprise Edition can handle twice as many clients as Windows 2000 Advanced Server.

Hardware utilization for the NetBench tests was measured using the counters listed in Figure 4. The results in Figure 8 show that Windows Server 2003 is highly optimized as a networking and file server compared to the previous Windows versions, delivering better performance while consuming less CPU power. Windows Server 2003 is also able to manage all available memory more efficiently and deliver the same amount of work with substantially fewer page faults. Furthermore, Windows Server 2003 can utilize all the hard drive power that the system offers to serve the file I/O requests from its clients.

Evolving to boost performance potential

The benchmark findings in this article have shown evolutionary performance improvements across the operating systems tested. This trend suggests that upcoming generations of operating systems will continue to make performance gains, using hardware resources more efficiently to finish tasks in less time. In the case of file service, response time for Windows Server 2003, Enterprise Edition was almost twice as fast as Windows 2000 Advanced Server, which translates into almost twice the number of requests that can be served in the same amount of time.

Microsoft optimizes its operating system kernel and its components to gain computation efficiency for both Web services and

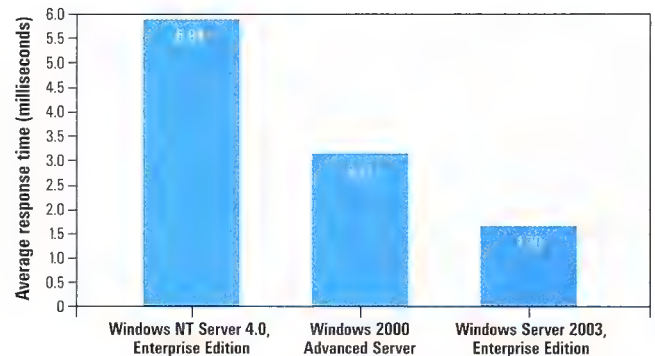


Figure 7. NetBench average response time results

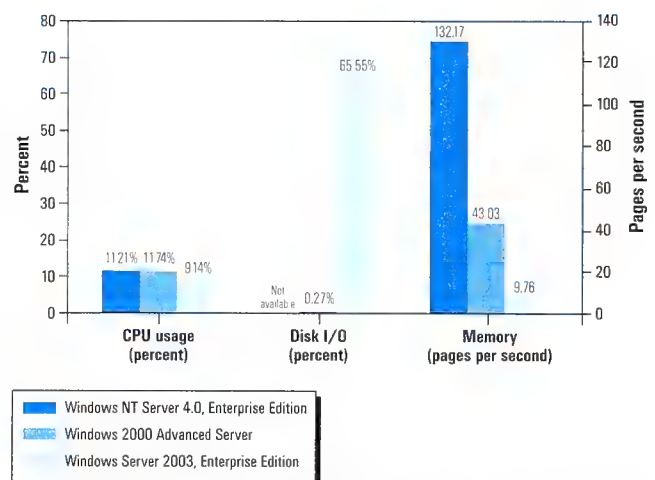


Figure 8. Hardware utilization during NetBench tests

network file sharing. Disk I/O was utilized hundreds of times more in Windows Server 2003 than it was in Windows 2000 Advanced Server, signifying the ability of Windows Server 2003, Enterprise Edition to use the right resource for the right job.

Users can easily mix and match different server applications running on the same physical server, each application maximizing the use of a different hardware component. Hardware utilization numbers also suggest that if both heavy-duty computation and server consolidation are required, an enterprise should employ servers such as the Intel Xeon processor-based Dell PowerEdge 4600, which incorporates the latest hardware technology and operating system.

Quynh Q. Nguyen (quynh_nguyen@dell.com) is a lead systems engineer consultant in the Server Operating Systems Engineering Department in the Dell Product Group—Enterprise Software Development. Quynh has a B.S. in Computer Science from Winona State University.

Min-John Lee (min-john_lee@dell.com) is a software engineering advisor in the Server Operating Systems Engineering Department in the Dell Product Group—Enterprise Software Development. He has an M.S. in Electrical and Computer Engineering from Northwestern University.

Using PCI Hot Plug

on Dell PowerEdge Servers Running Windows Server 2003

PCI (Peripheral Component Interconnect) Hot Plug is an industry-standard specification that can minimize server downtime by enabling a PCI device to be added or replaced while a server's operating system (OS) is running. Certain Dell™ PowerEdge™ server models enable many PCI devices to be hot-plugged while the Microsoft® Windows® Server 2003 OS is running.

BY SCOTT M. CALLAWAY AND FAISAL AHMED

More than ever, business environments require high availability of servers running business applications or providing storage. IT departments are often expected to keep uptime at the 99.999 percent industry benchmark.

PCI (Peripheral Component Interconnect) Hot Plug is an industry-standard specification that allows administrators to add or replace PCI devices while a server's operating system (OS) is running. Microsoft® Windows® Server 2003 supports PCI Hot Plug through the Advanced Configuration and Power Interface (ACPI). Certain Dell™ PowerEdge™ server models enable many PCI devices to be hot-plugged while the Microsoft Windows Server 2003 OS is running. As Windows Server 2003 is deployed on Dell servers, PCI Hot Plug will provide one way to enhance Dell high-availability implementations.

This article provides a basic overview of PCI Hot Plug functionality in Dell PowerEdge servers running Windows Server 2003, including required hardware configurations,

supported PCI devices, and best practices. PCI Hot Plug does not refer to failover functionality that allows automatic switching to a redundant device in the event of a failure. PCI Hot Plug describes the manual insertion and removal of PCI devices in a running server, as discussed in this article.

Supporting PCI Hot Plug using ACPI

Dell supports PCI Hot Plug functionality on servers that are deployed in environments where high availability is critical—currently, the PowerEdge 4400, 4600, 6300, 6350, 6400, 6450, 6600, 6650, and 8450¹ servers support this functionality, with additional hot-plug capable servers in development. Although PCI Hot Plug is mainly a single-server feature, it complements Dell clustering solutions by making individual server nodes more highly available. For example, a failed network interface card (NIC) used in a teaming or failover configuration can be replaced using PCI Hot Plug, and the end user will experience little or no system downtime.

¹ For PowerEdge 8450 servers, administrators must download and install the PHP filter driver for Windows 2000 Server and Windows Server 2003, and also enable the hot-plug controller (HPC) setting in the system's BIOS setup. A driver is not necessary for the other PowerEdge servers listed because hot-plug functionality is achieved through ACPI BIOS and the OS kernel for Windows 2000 Server and Windows Server 2003 only.

SUPPORTED PCI CARDS, MINIMUM REQUIREMENTS, AND SYSTEM CONFIGURATIONS

The following PCI adapters are supported by Windows Server 2003 and can be hot-plugged safely. Additional devices will be supported as they become available.

PowerEdge Expandable RAID Controllers (PERC):

- PERC 2/Single Channel (SC)
- PERC 2/Dual Channel (DC)
- PERC 2/Quad Channel (QC)
- PERC 3/SC
- PERC 3/DC
- PERC 3/Dual Channel Lite (DCL)
- PERC 3/QC
- PERC 4/SC
- PERC 4/DC

Network adapter cards:

- Intel® PRO/100 (82558)
- Intel PRO/100+ Server (82559)
- Intel PRO/100 S (82559)
- Intel PRO/100 S (82550)
- Intel PRO/1000
- Intel PRO/1000 F
- Intel PRO/1000 T
- Intel PRO/1000 XT
- Intel PRO/1000 XF
- Broadcom® 5700, 5701, 5703

Host bus adapters (HBAs):

- Emulex® LP 8000 Series
- Emulex LP 9000 Series
- QLogic® 2200, 2300, 2310, 2312, 2340

SCSI/IDE:

- LSI Logic™ CERC ATA100/4CH
- LSI Ultra320 SCSI Adapter Non-RAID
- Adaptec® SCSI card-39160 (Family Manager Set)
- Adaptec AHA-2940U2W, AHA-2944UW, SCSI Channel on Adaptec AHA-3944 PCI SCSI Controller (Family Manager Set)

Administrators should check that they have the latest system BIOS, firmware, and complex programmable logic device (CPLD) installed, especially with servers that are two or more years old. The Dell Support Web site, <http://support.dell.com>, offers this software for download.

In all servers that support PCI Hot Plug, a hardware component called the hot-plug controller (HPC) manages the power to the PCI slots and the connection to the host bus. The HPC detects power faults on the hot-plug expansion slots and if a fault occurs, power is removed from the slot immediately.

Operating systems that support PCI Hot Plug integrate the HPC in different ways. Windows 2000 Server and Windows Server 2003 support industry-standard ACPI BIOS (see Figure 1). ACPI defines standard interfaces for several configuration and power management features, including Hyper-Threading, hibernation, and standby. PCI Hot Plug uses ACPI to support the dynamic configuration of adapters added at run time.

The code needed to support the Dell HPC resides in the system ROM BIOS. When Windows Server 2003 loads, it notifies the system BIOS that it is an ACPI-compatible OS and scans the ROM looking for ACPI code. When the Windows Server 2003 OS finds the code for the HPC, it saves the machine instructions. When a hot-plug event occurs, Windows Server 2003 executes the ACPI code and hot-plug operations begin.

Note: On Windows 2000 Server and Windows Server 2003, the ACPI BIOS communicates with the HPC using the Intelligent Platform Management Interface (IPMI). However, Microsoft Windows NT® and Novell® NetWare® operating systems support PCI Hot Plug without using ACPI. On Windows NT and NetWare platforms, IPMI enables the HPC to communicate with PowerEdge Embedded Server Management (ESM) firmware. The firmware, in turn, talks to Windows NT and NetWare operating systems through the hardware

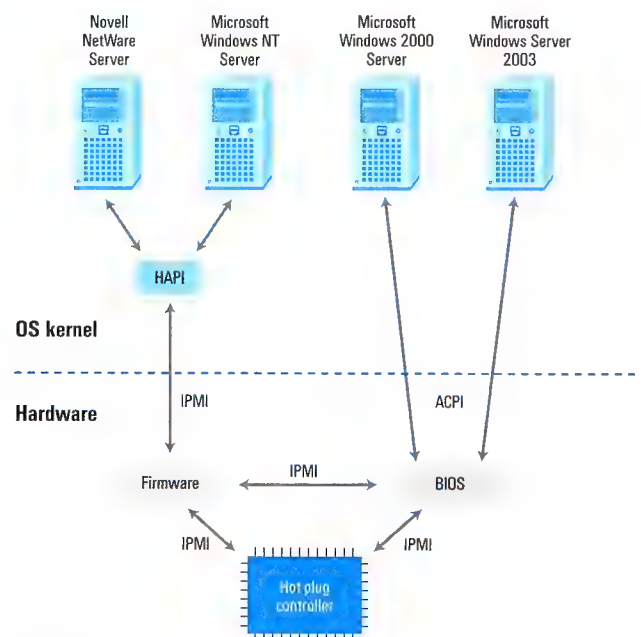


Figure 1. ACPI support of PCI Hot Plug



Yes! Sign me up for a Free Subscription to Dell **Power Solutions**, the magazine for direct enterprise solutions.

To request a free subscription or update your mailing address, please visit us online at www.dell.com/powersolutions_subscribe
or fax this card to +1 512-283-0109.

Check one: ☐ New subscription ☐ Address change ☐ Cancel subscription ☐ Renew subscription



Name: _____ Title: _____

Company/Agency: _____

Street Address: _____

City: _____ State/Province: _____

Country: _____ Zip/Postal Code: _____

Phone Number: _____ E-mail Address: _____



Please complete the questions listed below to receive your free subscription to Dell **Power Solutions**.
You must answer all questions to qualify for a free subscription.

1. Which of the following best describes your job function?

- | | |
|---|--|
| <input type="checkbox"/> Management (CFO, VP, Director) | <input type="checkbox"/> System administrator |
| <input type="checkbox"/> IT manager/supervisor | <input type="checkbox"/> Network administrator |
| <input type="checkbox"/> System analyst | <input type="checkbox"/> Project manager |
| <input type="checkbox"/> System engineer/technical consultant | <input type="checkbox"/> Marketing and sales |

2. What is the annual sales volume of your company?

- | | |
|--|---|
| <input type="checkbox"/> Less than \$5 million | <input type="checkbox"/> \$50-\$99 million |
| <input type="checkbox"/> \$5-\$9 million | <input type="checkbox"/> Greater than \$100 million |
| <input type="checkbox"/> \$10-\$49 million | <input type="checkbox"/> Greater than \$1 billion |

3. How large is your company in terms of employees?

- | | |
|--|---|
| <input type="checkbox"/> Less than 200 | <input type="checkbox"/> 1,001-5,000 |
| <input type="checkbox"/> 201-500 | <input type="checkbox"/> Greater than 5,000 |
| <input type="checkbox"/> 501-1,000 | |

4. What principal industry does your company serve?

- | | | |
|---|---------------------------------------|------------------------------------|
| <input type="checkbox"/> Manufacturing | <input type="checkbox"/> Education | <input type="checkbox"/> Retail |
| <input type="checkbox"/> Financial services and banking | <input type="checkbox"/> Healthcare | <input type="checkbox"/> Utilities |
| <input type="checkbox"/> Telecommunications | <input type="checkbox"/> Hospitality | <input type="checkbox"/> Other |
| <input type="checkbox"/> Government and military | <input type="checkbox"/> Internet/ISP | |

5. What Dell products does your company currently use?

- | | |
|---|--|
| <input type="checkbox"/> Desktops and notebooks | <input type="checkbox"/> Servers and storage |
| <input type="checkbox"/> All of the above | <input type="checkbox"/> None of the above |

6. What platforms are used in your company?

- | | |
|----------------------------------|--------------------------------------|
| <input type="checkbox"/> Windows | <input type="checkbox"/> UNIX |
| <input type="checkbox"/> Linux | <input type="checkbox"/> Novell |
| <input type="checkbox"/> Mixed | <input type="checkbox"/> Other _____ |

05/03



Yes! Sign me up for a Free Subscription to Dell **Power Solutions**, the magazine for direct enterprise solutions.

To request a free subscription or update your mailing address, please visit us online at www.dell.com/powersolutions_subscribe
or fax this card to +1 512-283-0109.

Check one: ☐ New subscription ☐ Address change ☐ Cancel subscription ☐ Renew subscription



Name: _____ Title: _____

Company/Agency: _____

Street Address: _____

City: _____ State/Province: _____

Country: _____ Zip/Postal Code: _____

Phone Number: _____ E-mail Address: _____



Please complete the questions listed below to receive your free subscription to Dell **Power Solutions**.
You must answer all questions to qualify for a free subscription.

1. Which of the following best describes your job function?

- | | |
|---|--|
| <input type="checkbox"/> Management (CFO, VP, Director) | <input type="checkbox"/> System administrator |
| <input type="checkbox"/> IT manager/supervisor | <input type="checkbox"/> Network administrator |
| <input type="checkbox"/> System analyst | <input type="checkbox"/> Project manager |
| <input type="checkbox"/> System engineer/technical consultant | <input type="checkbox"/> Marketing and sales |

2. What is the annual sales volume of your company?

- | | |
|--|---|
| <input type="checkbox"/> Less than \$5 million | <input type="checkbox"/> \$50-\$99 million |
| <input type="checkbox"/> \$5-\$9 million | <input type="checkbox"/> Greater than \$100 million |
| <input type="checkbox"/> \$10-\$49 million | <input type="checkbox"/> Greater than \$1 billion |

3. How large is your company in terms of employees?

- | | |
|--|---|
| <input type="checkbox"/> Less than 200 | <input type="checkbox"/> 1,001-5,000 |
| <input type="checkbox"/> 201-500 | <input type="checkbox"/> Greater than 5,000 |
| <input type="checkbox"/> 501-1,000 | |

4. What principal industry does your company serve?

- | | | |
|---|---------------------------------------|------------------------------------|
| <input type="checkbox"/> Manufacturing | <input type="checkbox"/> Education | <input type="checkbox"/> Retail |
| <input type="checkbox"/> Financial services and banking | <input type="checkbox"/> Healthcare | <input type="checkbox"/> Utilities |
| <input type="checkbox"/> Telecommunications | <input type="checkbox"/> Hospitality | <input type="checkbox"/> Other |
| <input type="checkbox"/> Government and military | <input type="checkbox"/> Internet/ISP | |

5. What Dell products does your company currently use?

- | | |
|---|--|
| <input type="checkbox"/> Desktops and notebooks | <input type="checkbox"/> Servers and storage |
| <input type="checkbox"/> All of the above | <input type="checkbox"/> None of the above |

6. What platforms are used in your company?

- | | |
|----------------------------------|--------------------------------------|
| <input type="checkbox"/> Windows | <input type="checkbox"/> UNIX |
| <input type="checkbox"/> Linux | <input type="checkbox"/> Novell |
| <input type="checkbox"/> Mixed | <input type="checkbox"/> Other _____ |

05/03

Apply
Postage
Here

DELL COMPUTER CORPORATION
POWER SOLUTIONS
ONE DELL WAY, RR #5, BOX **8503**
ROUND ROCK, TX **78682**
USA

Apply
Postage
Here

DELL COMPUTER CORPORATION
POWER SOLUTIONS
ONE DELL WAY, RR #5, BOX **8503**
ROUND ROCK, TX **78682**
USA

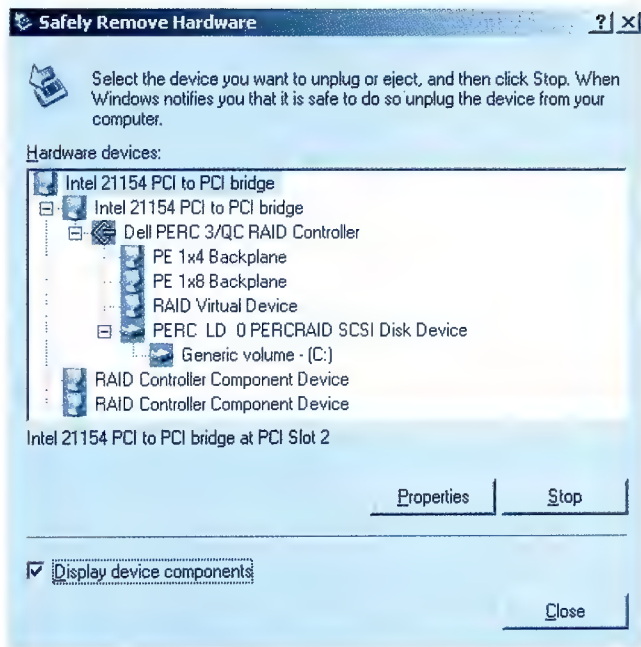


Figure 2. The Safely Remove Hardware application

application programming interface (HAPI) using hot-plug software provided by Dell.

Understanding PCI Hot Plug functionality

PCI Hot Plug functionality falls into two major categories: hot add, in which an adapter is added to the system; and hot replace, in which an adapter in the system that is no longer operating correctly is replaced by another adapter.

For example, PCI Hot Plug support enables system administrators to add I/O adapters while the system is still running. System administrators who need to add users or balance network traffic can add an extra LAN adapter to a Dell PowerEdge server without powering down the server. The device driver can then be loaded and the service started while users continue to work uninterrupted.

After the addition or replacement of a device, the slot must be powered on using the slot power button next to each slot,² reset, configured, and brought online in an orderly manner—the same steps that occur during a traditional system power-up sequence. This process involves loading the device driver for the adapter and starting the appropriate service that uses the adapter.

Safely Remove Hardware application helps prevent data loss

The failure of an adapter that the OS uses can lead to data loss or data-access problems. For example, if the disk drive adapter

fails to respond to device driver requests that were initiated by the OS, the user might not be able to retrieve the desired data. Because PCI Hot Plug enables system administrators to replace a failed adapter while the system continues to operate, data loss is prevented and users can retrieve the data after the new adapter is replaced and initialized.

To safely remove a PCI device without incurring data loss, administrators should use the Safely Remove Hardware application if its icon appears in the Windows Server 2003 notification area. When a hot-plug PCI device needs to be removed, system administrators can double-click this icon, displaying the Safely Remove Hardware dialog box (see Figure 2). This application informs the Windows OS of the imminent removal of the PCI hot-pluggable device, enabling Windows to sequentially stop processes and services related to the device's removal.

The removal of a PCI device without using the Safely Remove Hardware application is known as a surprise removal. Because surprise removals do not notify the operating system of the device's removal, unexpected results may ensue. For example, if a Dell PowerEdge Expandable RAID Controller 3, Quad Channel (PERC 3/QC) is removed and write-caching is enabled, data loss could occur. (Although the default installation of Windows Server 2003 disables write-caching, a RAID controller may have enabled it for performance reasons in certain server configurations.) To prevent such problems, Dell and Microsoft recommend using the Safely Remove Hardware application to remove hot-plug devices.

Hot-replace capability improves failover support

Although totally uninterrupted service from a server cannot be guaranteed, PCI Hot Plug allows system administrators to replace an adapter that has stopped operating or appears to have a problem—without bringing down the server. When the hot-plug

services module that is built into the OS kernel recognizes that an adapter has failed, the module attempts to unload the device driver, and detaches the device from the services that use the adapter. Services are suspended until administrators use the Safely Remove Hardware application to replace the failed adapter and bring the new one online.

Some OS services cannot sustain the loss of certain adapters. In an effort to move beyond this limitation, Dell works closely with

As Windows Server 2003

is deployed on Dell

PowerEdge servers,

PCI Hot Plug will provide

one way to enhance

Dell high-availability

implementations.

² For the power button layout of an individual server, please refer to the server's system documentation.

other vendors to provide a failover solution for selected adapters that allow the OS to continue to operate without interruption. To support the failover capability of an adapter, a redundant pair of adapters is sometimes used. Two identical adapters (a primary and a backup) are inserted in the I/O expansion slots. A


PowerEdge server	Green LED action	Amber LED action	Safe to remove card?	Condition
PowerEdge 4400 and 64xx	Off	Off	Yes	Slot power is off
	On	Off	No	Slot power is on
	Flash	Off	No	Slot power is being powered up or powered down
	Off	Flash slow	Yes	Power-up fault
	On	Flash fast	No	User application or device driver identifying a slot
PowerEdge 4600 and 66xx	On	Off	No	Slot power is on
	Off	Flash slow	Yes	Slot power is off; attention needed
	Off	Flash fast	No	Slot power is on; attention needed
PowerEdge 63xx	On	Off	No	Slot power is on; no error
	Off	Flash slow	Yes	Slot power is off; startup default; adapter has shorted or is incorrectly installed; attention needed
	Off	Flash fast	No	Slot power is on; adapter or software fault; attention needed
	Off	Off	Yes	Slot power is turned off or adapter is installed in slot
PowerEdge 8450	Off	Off	Yes	Slot power is off
	On	Off	No	Slot power is on; no error
	Off	On	Yes	Slot power is off; power-related expansion-card connection error
	On	On	No	Slot power is on; expansion card or software driver fault
	Flash	Off	No	Slot is being powered up or powered down

Figure 3. Slot LED conditions for various PowerEdge servers

special device driver closely monitors the activities of the primary adapter. If the primary adapter fails, the device driver quickly switches the operations to the backup adapter, and the system administrator is notified. The administrator can then hot-replace the failed adapter and bring it back online, thus restoring the redundant failover capability.

The LED table in Figure 3 indicates the state of a PCI device in a hot-plug slot.

Enhancing the availability of systems using PCI Hot Plug technology

In today's business environments, servers running business applications and storage must maintain standards of high availability. For enterprises that require a highly available server with a flexible and expandable I/O subsystem that can be modified while the system is running, a hot-plug capable Dell PowerEdge server running Windows Server 2003 is an excellent choice to increase the availability of the IT infrastructure. 

Scott M. Callaway (scott_m_callaway@dell.com) is a software engineer in the Server Operating Systems Engineering Department in the Dell Product Group. Scott spent his first few years at Dell on the team responsible for the Dell IT department's early adoption of Windows 2000 as its corporate desktop and server standard and is now working on Dell server products that will support Windows Server 2003. Scott has a B.S. in Management from Stephen F. Austin University and more than 10 years of industry experience in Microsoft operating systems.

Faisal Ahmed (faisal_x_ahmed@dell.com) is a software engineer in the Server Operating Systems Engineering Department in the Dell Product Group. Faisal has been working to add hardware platform support for Windows Server 2003. His responsibilities include developing PCI Hot Plug software for Dell PowerEdge servers. He has a degree in Computer Systems Engineering from the University of Houston at Clear Lake.

FOR MORE INFORMATION

<http://www.microsoft.com>

<http://support.dell.com>

PCI Hot Plug allows system administrators to replace an adapter that has stopped operating or appears to have a problem—without bringing down the server.

High-Availability Clustering Capabilities

in Windows Server 2003

Cluster environments enable multiple systems to provide higher availability of applications and services than a single system can provide. With the release of the Microsoft® Windows® Server 2003 operating system, Microsoft enhances the availability, scalability, and manageability of its Windows clustering products. Dell™ PowerEdge™ Clusters are poised to take advantage of these new features and enhancements. This article discusses several new clustering features and enhancements in Windows Server 2003.

BY MIKE KOSACEK, NAM NGUYEN, AND MARK TIBBS

As system uptime becomes increasingly important for business-critical and mission-critical operations, many enterprise organizations are turning to server cluster configurations to handle high workloads and availability requirements. Clustering combines multiple systems and enables them to work together, resulting in higher availability, scalability, and manageability. Clustering solutions also can enable consolidation of servers, storage, and IT personnel across an organization.

Standards-based high-availability clustering solutions from Dell include proven, tested, integrated, and reliable products that help ensure the availability of critical applications and services at all times. Clusters of Dell™ PowerEdge™ servers running the Microsoft® Windows® Server 2003 operating system are well suited

for the enterprise and can provide the following enhancements:

Higher availability. Failover clusters help achieve higher availability. When a software or hardware failure occurs on one node, applications or services on that node are moved, or failed over, to another node in the cluster, thereby minimizing downtime. The applications or services are then restarted on this node. From the client's perspective, only a slight delay or decrease in performance may be noticed.

Greater scalability. Microsoft clustering technologies now enable greater scale-out capabilities by allowing up to eight nodes per cluster. An increased number of nodes allows more efficient use of storage systems and storage area networks (SANs). Greater scale-up capabilities can

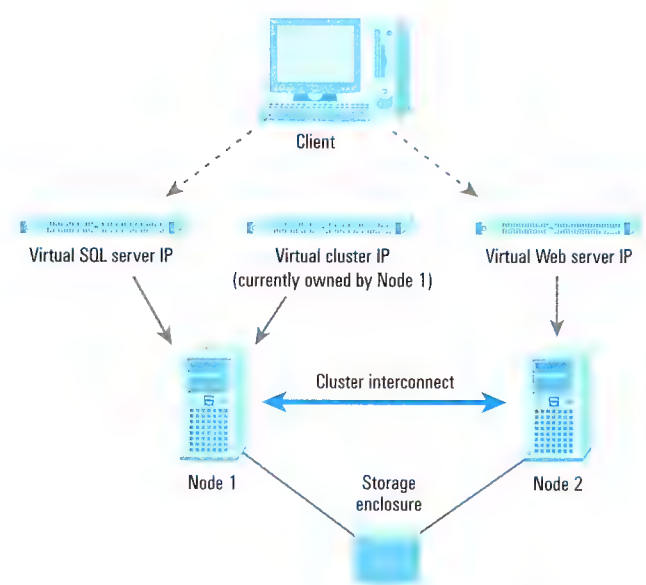


Figure 1. An example of an active/active cluster using Microsoft Cluster Service

be achieved through support in Windows Server 2003 for symmetric multiprocessing (SMP) with more memory and a larger number of processors—up to 64 CPUs per node (for Intel® Itanium® processor-based servers). Load-balancing application instances across several computer systems also can facilitate scalability. If a node in a cluster fails, the cluster can redistribute the workload across the remaining members of the cluster.

Better manageability. Managing a cluster of nodes, rather than managing each computer system, can provide simpler administration. In addition, virtual servers ease management by providing a single point of administration for the cluster resources. Client applications connect to the virtual servers, allowing the cluster-aware application to reconnect without manual intervention.

Achieving high availability with Microsoft Cluster Service

Microsoft Cluster Service (MSCS) is a component of the Microsoft Windows 2000 Advanced Server; Windows 2000 Datacenter Server;

Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition operating systems. The MSCS architecture allows for a virtual server to fail over between multiple physical servers running the cluster service. Clients connect to the IP address of the virtual server (see Figure 1), and are then directed to the physical server node that currently owns the clustered application or service. The client is unaware of which physical server it is actually using. If the physical server fails, the virtual server (and associated resources) fail over to another node and the client maintains its connection through the virtual IP.

Multiple virtual servers may exist in a single cluster. In an active/active or activeⁿ cluster (where *n* equals the number of active nodes), each node in the cluster owns at least one virtual server. In an active/passive or *N+1* cluster (where *N* and *I* are the numbers of active and passive nodes, respectively), at least one cluster node does not own any virtual server and functions as the backup for the active cluster nodes.

Although an MSCS cluster can support multiple applications or multiple instances of a single application, each application or application instance accesses a separate data set. An MSCS cluster provides application failover, but does not load balance a single application instance across multiple servers.

An application may fail over automatically if a hardware or software failure occurs on one node of the cluster. Administrators also can initiate failover manually, often to allow routine maintenance to be performed on the cluster nodes. Manual failover can be useful when installing software updates or performing rolling upgrades, a process in which one node is upgraded while other nodes continue to host applications.

Enhancing clustering capabilities in Windows Server 2003

Many new features and enhancements are available in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition. MSCS is now an integrated part of the operating system, no longer an optional component. The following sections present more detail on the new features that affect MSCS (see Figure 2 for a comparison of the product features).

Windows Server 2003 product family	Number of CPUs	Maximum RAM	Number of server cluster nodes supported	Number of Network Load Balancing (NLB) cluster nodes supported	Available in 64-bit version?
Web Edition	1 or 2	2 GB	N/A	32	No
Standard Edition	1 to 4	4 GB	N/A	32	No
Enterprise Edition	1 to 8	32 GB (x86 processor) 64 GB (Intel Itanium processor)	8	32	Yes
Datacenter Edition	1 to 64	64 GB (x86) 512 GB (Intel Itanium processor)	8	32	Yes

Figure 2. Windows Server 2003 product family comparison

Easier cluster setup and configuration

Before creating a cluster, administrators can analyze the hardware and software configurations and identify potential compatibility problems. Administrators can set up, reconfigure, and remove a cluster configuration without the installation CD and without rebooting the nodes. Cluster setup can be fully or partially automated

using command-line (cluster.exe) batch files or scripts. Multiple nodes can be added to the cluster by performing a single operation from a cluster node or a remote management station. To uninstall a cluster node, administrators simply evict the node from the cluster.

Administrators can now delete cluster resources without taking them offline first. The MSCS account password can be changed on multiple nodes simultaneously without taking any nodes offline, resulting in higher availability for the cluster.

Clusters of Dell

PowerEdge servers

running the Microsoft

Windows Server 2003

operating system

are well suited for

the enterprise.

Server. A greater number of nodes enables the cluster to tolerate multiple server failures; two-node configurations can tolerate only one node failure. Larger cluster sizes give administrators more options to deploy applications and configure failover policies that better match business needs. Figure 2 shows the number of nodes supported in each edition.

Different applications or multiple instances of the same application (each with separate data sets) can be deployed across multiple nodes. MSCS provides group affinity support that allows applications to fail over to spare nodes before active nodes. As a result, $N+1$ configurations are easily supported, enabling deployment of larger clusters of Microsoft Exchange Server or consolidation of file-and-print servers.

New MNS quorum

The majority node set (MNS) quorum is an optional cluster resource for the quorum device. This resource allows the quorum to reside on the local disks of the cluster nodes, eliminating the need for a physical disk on the shared storage, which could be a single point of failure in the cluster. This feature eases the tasks for deploying geographically dispersed clusters and enables administrators to deploy low-cost clusters without the shared storage. Such clusters are also well suited to host services or applications that do not require storage or data to be shared among the nodes.

Larger default quorum size

The default size for the quorum log has been increased from 64 KB to 4096 KB to accommodate large numbers of file or printer shares. Administrators no longer need to specify the quorum disk during the installation; the smallest disk with an NT file system (NTFS) volume of more than 50 MB (Dell recommends a 1 GB disk) will be selected for the quorum by default. After the cluster has been configured, the quorum can be moved to another disk, if desired.

Simpler cluster resource setup

Setting up clustered printers and Microsoft Distributed Transaction Coordinator (MS DTC) is much easier in Windows Server 2003. The Print Spooler or the MS DTC resource can be configured on one node and then automatically replicated to the remaining nodes. In addition to creating resource dynamic-link libraries (DLLs), which are usually written in C/C++, programmers can use VBScript or JScript® development software to write resource plug-ins for existing applications. Making existing applications cluster-aware allows the cluster service to monitor and manage the applications, easing administration.

More cluster nodes

Windows Server 2003, Enterprise Edition and Datacenter Edition, supports a maximum of eight cluster nodes—up from two in Windows 2000 Advanced Server or four in Windows 2000 Datacenter

Greater maximum supported physical and addressable memory

In the 64-bit versions of Windows Server 2003, the maximum supported physical memory has been increased from 8 GB in Windows 2000 Advanced Server, and 64 GB in Windows 2000 Datacenter Server, to a maximum of 512 GB. The maximum addressable memory space extends to 16 TB. Computing-intensive applications such as Microsoft SQL Server 2000 Enterprise Edition (64-bit) can take advantage of increased memory space in the cluster environment.

Full integration with Active Directory

MSCS can now achieve full integration with the Microsoft Active Directory® directory service. The cluster virtual server can be published in the Active Directory as a computer object. Being published as a computer object enables Kerberos authentication for services being hosted on a virtual server and allows cluster-aware and Active Directory-aware services to publish service provider information specific to the host virtual server.

New WMI support

In addition to the MSCS application programming interfaces (APIs), Windows Management Instrumentation (WMI) can be used to manage the clusters. Under the WMI environment, administrators can create, delete, start, or stop cluster resources. The cluster state

and the status of cluster applications also can be monitored, and any cluster state change can be propagated to applications that have subscribed to such WMI events.

New network enhancements

MSCS examines the state of the public network before arbitrating for the quorum disk. This examination reduces the performance hit from disk arbitrations and prevents a scenario in which the surviving node has a bad public network interface and thus cannot provide services to the clients. Media sense is disabled by default and administrators no longer must set the DisableDHCPMediaSense registry key to preserve the TCP/IP stack when the network connection is lost.

Multicast communication is automatically selected for heartbeats between the nodes in large cluster configurations if the network infrastructure can support it. With multicast support, the amount of network traffic is reduced, which can be very beneficial in large or geographically dispersed clusters. If multicast communication fails, the heartbeat will automatically revert to unicast.

More flexible file system support

Previously, the number of shared disks per cluster was limited to 23 (drive letters D through Z). New support for volume mount points on the cluster's shared disks eliminates that limitation. Client-side caching (CSC) allows the clients to cache data on a clustered share down to the clients' local disks. When the file is closed, the data from the local disks is copied back to the clustered share, making failover transparent to the clients.

The Distributed File System (DFS) has several improvements. Multiple stand-alone roots can provide more flexibility in deploying the DFS namespace. Each DFS root can fail over individually, resulting in improved failover time. In active/active configurations, multiple stand-alone roots can run actively on multiple cluster nodes. The Encrypting File System (EFS) also is supported on the clustered file shares to further enhance the file system security.

If the storage hardware supports dynamic expansion of a disk,

that disk volume can now be extended while it is online using the included Diskpart.exe utility. This utility allows more disk space to be added to the cluster without interrupting its applications and services.

Better SAN support

Windows Server 2003 can support booting from the SAN, allowing the system disk with the page file and dump file to be on the

same external storage system as the cluster's shared disks. Booting from the SAN permits centralized backups and simpler maintenance. Also new, the targeted device reset mechanism minimizes the possibility of disruption from a bus reset during cluster arbitration for other devices attached to the same host or to other hosts sharing the same storage system(s).

Easier cluster backup and restore

Backup and restoration of cluster configurations have become much easier. The Windows Server 2003 operating system backup utility (NTBackup.exe) can now back up the cluster configuration and restore that configuration locally or to all the cluster nodes. With Automated System Recovery (ASR), a cluster node can be restored if its local cluster database is damaged, or if its system disk fails or becomes corrupted. The ASR feature also can help restore the shared disks, including the disk signatures.

Quicker, easier troubleshooting and failure recovery

Cluster failures can be diagnosed faster and more easily in Windows Server 2003 because of various error logging enhancements. More failure reason codes have been added to help determine whether an application has failed or been taken offline by an administrator. The system event log shows both cluster resources that have been successfully failed over and cluster resource failures.

A new software tracing feature allows administrators to troubleshoot cluster issues without bringing the node down and loading checked build (debug) versions of the DLLs. A separate setup log (%SystemRoot%\System32\Logfiles\Cluster\CICfgSrv.log) is created during the cluster setup to help diagnose problems during the installation. When chkdsk is run against a cluster's shared disk, a chkdsk log is created to assist administrators in troubleshooting and fixing problems.

With the new Error Levels (info, warn, err), administrators can easily find the entries of interest without sifting through the cluster logs. The Local Server Time Stamp in the cluster log can help correlate the cluster log with the system event logs. A cluster object file (%windir%\Cluster\Cluster.obj) with mappings between globally unique identifiers (GUIDs) and cluster resource names is also automatically created and updated.

A new diagnostics tool (ClusDiag.exe) in the Windows Server 2003 Resource Kit allows cluster logs and event logs from all nodes in the cluster to be correlated. This diagnostic tool also can be used to validate the cluster configuration by running stress tests against

Windows Server 2003

can support booting

from the SAN,

allowing centralized

backups and simpler

maintenance.

Cluster failures can be
diagnosed faster and
more easily in Windows
Server 2003 because
of various error logging
enhancements.

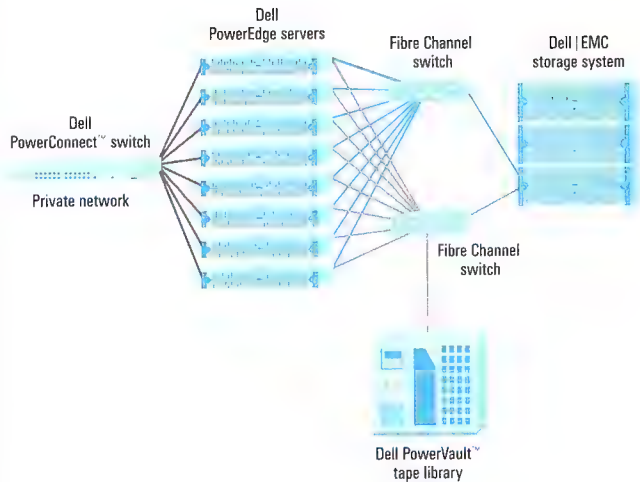


Figure 3. Windows Server 2003 eight-node cluster

the server nodes, the storage, and the cluster infrastructure before putting the configuration into a production environment.

Exploring Dell clustering products

Dell offers highly available two-node clusters based on SCSI storage and running Windows 2000 Advanced Server or Windows Server 2003, Enterprise Edition. The SCSI-based clusters are well suited for file-and-print servers, messaging servers, or database servers for smaller workgroups.


Dell also offers two-node to eight-node clusters based on Dell | EMC Fibre Channel storage systems. Dell two-node clusters support Windows 2000 Advanced Server, and a four-node cluster is offered with Windows 2000 Datacenter Server. For Windows

Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, Dell supports up to eight nodes in a failover cluster (see Figure 3). Windows Server 2003 clusters can be installed with a fresh installation or with a rolling upgrade from a Windows 2000 cluster.

Each cluster component in the Fibre Channel clusters is fully redundant to provide high availability. Any component in the cluster can be repaired or replaced and the cluster can be upgraded with the latest software, driver, and firmware while the cluster is still online serving its clients.

Cluster components are scalable to meet enterprise requirements. The host bus adapters (HBAs), switches, storage systems, Fibre Channel disks, and tape libraries can operate at 2 Gbps. Multiple clusters and nonclustered servers can be connected to redundant fabrics and can share multiple storage systems and tape libraries. The storage systems offer optional advanced features such as storage mirroring or snapshot copies of disks for backup. The Dell Fibre Channel cluster solutions are well suited for enterprises that require high availability on messaging servers, database servers with large data sets, and consolidation of applications, servers, or storage.

Increasing uptime for critical applications

The new features of Windows Server 2003 help enable easier deployment, increase performance, enhance security, and promote increased uptime for clustered applications and services. Dell PowerEdge clusters running Microsoft Windows operating systems combine the high performance, reliability, and availability features of Dell server and storage products with the powerful clustering capabilities of MSCS to help enterprises achieve high availability on critical applications. 

Mike Kosacek (michael_kosacek@dell.com) is a senior member of the High-Availability Cluster Development Group at Dell. As lead engineer on several Dell products, Mike develops, tests, integrates, and certifies cluster solutions for the Dell server and storage product lines. Mike has an Electronics Technology degree and is a Microsoft Certified Systems Engineer (MCSE).

Nam Nguyen (nam_nguyen@dell.com) is a systems engineer in the High-Availability Cluster Development Group at Dell. Nam is the lead engineer for several Dell Fibre Channel PowerEdge Cluster products. He has both a B.S. and an M.S. in Electrical Engineering from The University of Texas at Austin.

Mark Tibbs (mark_tibbs@dell.com) is a member of the High-Availability Cluster Development Group at Dell. Mark was an officer in the Air Force. Tours included the Middle East and Central America. He also led a team that managed the backbone networks for the Pentagon. Mark has a B.S. in Computer Information Systems from the University of Tampa and an M.S. in Computer Systems Management from the University of Maryland.

FOR MORE INFORMATION

Dell High-Availability Clusters:
<http://www.dell.com/clusters>

Windows 2000 clustering technologies:
<http://www.microsoft.com/windows2000/technologies/clustering/default.asp>

Windows Server 2003 clustering technologies:
<http://www.microsoft.com/windows.netserver/evaluation/overview/technologies/clustering.msp>

Using Microsoft Cluster Service

to Enhance Dell Clusters

Server clusters offer high availability, scalability, and manageability for resources and applications. Microsoft[®] Cluster Service (MSCS) helps keep applications on a cluster running, even if a server fails. This article highlights how Cluster Administrator, available with the Microsoft Windows[®] Server 2003 operating system, facilitates cluster installation.

BY MAHMOUD AHMADIAN AND DANIEL MOGES

Clustered computer environments allow many servers to be managed as if they were one server. In addition, they scale easily and provide high availability. Clusters are becoming more common in enterprise computing infrastructures, and more manufacturers are offering products that enable cluster configurations. A typical Dell[™] high-availability cluster uses two or more Dell PowerEdge[™] servers, Dell PowerVault[™] storage (SCSI or Fibre Channel), Dell|EMC storage, and Dell PowerConnect[™] switches. Microsoft offers software, such as the Microsoft[®] Windows[®] Server 2003, Enterprise Edition operating system (OS) and the Microsoft Cluster Service (MSCS),¹ to provide maximum availability for mission-critical applications. Cluster-aware applications, including Microsoft SQL Server 2000, Microsoft Exchange Server 2003, and the Oracle[®] Fail Safe failover software, are designed to seamlessly integrate with MSCS.

MSCS: Supporting cluster functionality

Windows Server 2003 (both the 32-bit and 64-bit versions) enhances cluster functionality by supporting up to eight servers in a cluster. The OS includes the Majority Node Set (MNS) resource, which allows clusters to be geographically dispersed; a quorum resource, which maintains the configuration data necessary for cluster recovery; and low-cost, appliance-like solutions that do not use shared storage but use software methods such as log shipping and software disk replication. MNS resource configuration is outside the scope of this article.

The Windows Server 2003 Resource Kit also provides several cluster utilities that help to troubleshoot problems with server clusters. The ClusDiag tool, for example, helps perform triage and validate cluster configurations. Another tool that simplifies cluster backup and recovery is Cluster Tool.²

¹ Included only in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition.

² Microsoft Server 2003 Resource Kit.

Cluster Administrator

This article explains how administrators can use Cluster Administrator and its New Server Cluster wizard to maximize ease of use for Windows Server 2003.

Cluster Administrator: Creating clusters and administering resources

Cluster Administrator allows administrators to create cluster objects and displays information about cluster groups and resources through a single console. This one-to-many capability eases configuration, management, and monitoring of cluster installations. After a system is added as a member of a cluster object, Cluster Service is installed on that system. When using Cluster Administrator on systems running Windows Server 2003, administrators can let a wizard guide the creation of a new cluster.

The New Server Cluster wizard prompts for the cluster network name as well as the host name of the first system (node) in the cluster. The wizard then systematically analyzes the configuration on each node by following these steps:

1. **Check for existing cluster.** The wizard verifies that no network name conflict exists by checking Domain Name System (DNS) records and NetBIOS names. If this step of the test fails, so does the configuration; the name conflict must be resolved before proceeding to the next step.
2. **Establish node connection(s).** The wizard connects to the server that will become a member of the cluster object, and starts the cluster configuration server.
3. **Check node feasibility.** The wizard checks for cluster feasibility of the node by validating the node OS version.
4. **Find common resources on nodes.** The wizard collects managed resources and reports on each resource type found. A quorum-capable resource, either Local Quorum or MNS, is required. Next, the wizard looks for and reports disks on the same storage bus as the boot disk, because these cannot be managed as a cluster resource. If a sharable quorum resource is not found, a local quorum cluster will be created. The wizard also collects network resource information and reports any concerns. For example, if a network adapter is enabled for Dynamic Host Configuration Protocol (DHCP), the wizard generates a warning indicating that the node in question cannot be supported as a member of the cluster.
5. **Check cluster feasibility.** The wizard verifies cluster membership to ensure that the node is not already a member of

allows administrators
to create cluster objects
and displays information
about cluster groups
and resources through
a single console.

an existing cluster. It then compares the processor architecture of the node with that of the other nodes in the cluster to detect any interoperability issues. Next, the wizard performs a test to detect any drive letter collisions among the nodes in the cluster. Finally, the wizard checks that all nodes have access to a disk resource that could be used as a quorum resource.

New Server Cluster wizard: Ensuring accessibility of the quorum disk

Every node in a cluster must be able to access the quorum disk resource. The storage on which cluster disks reside is known as shared storage. Several scenarios exist in which administrators, to ensure access to the quorum, should change the quorum selection disk in the New Server Cluster wizard.

Non-shared disks residing on split backplane. Administrators should use care when a system that uses a split backplane configuration participates in a Windows Server 2003 cluster. In such systems, if both channels of the split backplane are connected to SCSI controllers within the system, two SCSI buses are formed. For purposes of this discussion, these buses can be called the primary bus and the secondary bus. The New Server Cluster wizard views the disks residing on the same SCSI bus as the system boot disk (the primary bus) as non-shared, and therefore does not consider them eligible for clustering. On the other hand, the wizard views disks on the secondary bus as shared and therefore eligible for clustering. However, in this case the wizard is wrong, since the secondary bus is still local to the system and is not accessible by other systems in the cluster. Administrators must correct this problem.

Quorum disk selection in Typical configuration. The Typical (full) configuration option within the New Server Cluster wizard allows all SCSI disks on the secondary bus to be configured as cluster disk resources. To avoid using a disk on the secondary bus as a quorum disk, the New Server Cluster wizard allows administrators to select the quorum disk. In Figure 1, the New Server Cluster wizard has chosen disk E as the quorum disk. In this scenario, disk E resides on the secondary bus, and is therefore non-shared. Because disk E is not accessible to all nodes in the cluster—that is, does not reside on the shared storage—administrators should change this selection, and select a disk that is accessible by all cluster nodes. To help ensure successful cluster operations, administrators must manually remove all disks not accessible by all nodes from the list of cluster disk resources once cluster configuration is complete.

The one-to-many
capability of Cluster
Administrator eases
configuration, manage-
ment, and monitoring
of cluster installations.

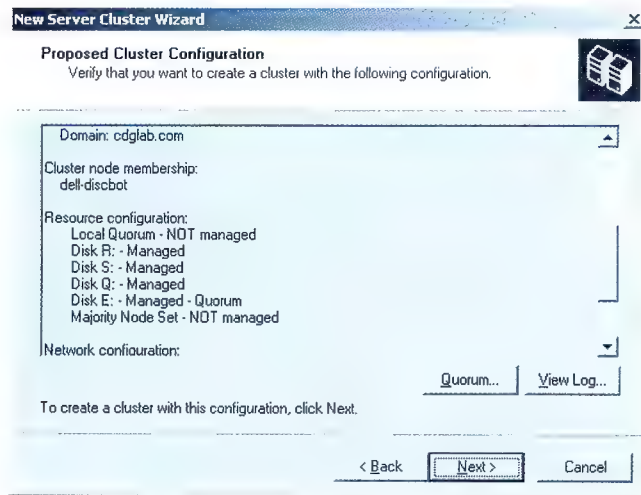


Figure 1. New Server Cluster wizard: Typical configuration option

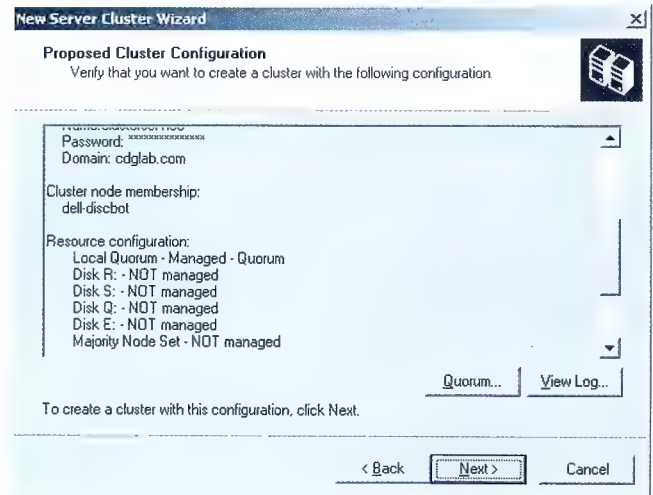


Figure 2. New Server Cluster wizard: Advanced configuration option

Quorum disk selection in Advanced configuration. When selecting the Advanced (minimum) configuration option, only one disk may be selected and configured as a cluster disk resource; this disk becomes the quorum disk (see Figure 2). In the Advanced configuration option, the local disk (boot disk) is selected as a quorum disk, and all other disk resources are seen as unmanaged disk resources.

Because a local quorum is not accessible by all nodes in the cluster, administrators should change the quorum disk selection from the boot disk to one of the disks accessible to all nodes in the cluster. After the setup is completed, administrators can add all other shared disk resources from Cluster Administrator by manually creating, for each disk on the shared storage, a new Disk Group and Physical Disk Resource.


Answer file: Allowing unattended cluster installation

To perform an unattended installation and configuration of Cluster Service on systems running Windows Server 2003, Enterprise Edition, administrators can add entries to the unattended installation answer file. This file, called Unattended.txt, is located on the OS CD. Entries under the cluster installation heading in this file are used by the installation program to configure clustering. In a cluster, at least two networks, one public and one private, must be configured. Additionally, Microsoft recommends that the associated network cards have static IP addresses. In the networking

The enhancements
to cluster installation
and configuration in
Windows Server 2003
ease deployment of
server clusters.

portion of the answer file, administrators must enter the appropriate information, including the network protocol and static IP address, to configure the networks for a cluster.

Windows Server 2003: Easing cluster deployment

The enhancements to cluster installation and configuration in Windows Server 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, ease deployment of server clusters. New features, such as support for up to eight nodes and support for Majority Node Set, provide for even better availability, whether clusters are local or geographically dispersed. Windows administrative tools, such as Cluster Administrator and the New Server Cluster wizard, also streamline cluster installation and management. By following the procedures outlined in this article, administrators can maximize their Windows Server 2003 cluster performance. 

Mahmoud Ahmadian (mahmoud_ahmadian@dell.com) is a development engineer in the High-Availability Cluster Development Group at Dell. He has an M.S. in Computer Science from the University of Houston.

Daniel Moges (daniel_moges@dell.com) is a systems engineer in the High-Availability Cluster Development Group at Dell. He has a B.S. in Electrical Engineering from the University of Illinois at Urbana-Champaign.

FOR MORE INFORMATION

http://www.dell.com
http://www.microsoft.com

Microsoft SQL Server 2000

Scalability:

The Rosetta Genomics 10 TB Human Genome Database

Microsoft® SQL Server™ 2000 Enterprise Edition running on an industry-standard Dell™ enterprise server supports very large database (VLDB) applications. This article discusses storage, loading, and configuration options used in a real-world application of Microsoft SQL Server 2000 Enterprise Edition in a 10 TB database application at the Rosetta Genomics™ company. Engineers at Microsoft labs studied the performance effects of several configuration alternatives for building and loading the Rosetta Genomics database, and this article presents the most effective solutions.

BY JOE BELSON AND HAO CHEN

More companies are looking for ways to maximize return on investment (ROI) for their IT departments, while striving to manage ever-increasing amounts of data. Microsoft® SQL Server™ 2000 Enterprise Edition provides very large database (VLDB) support with a much reduced total cost of ownership (TCO) compared to the mainframe or UNIX® installations that have been the predominant platforms for multiple-terabyte databases.

Genomic research is one field where the use of multiple-terabyte databases has become essential. Genomic data is a major source for discovery and development of new products by the pharmaceutical and biotechnology industries. However, the complexity and sheer size of genomic data make the task of genomic data mining and analysis a daunting one: genomic-sequence data is accumulating at an exponential growth rate, having doubled every 18 months in the past 5 years.

Until recently, genomic applications required mainframes or supercomputers that ran custom data management applications, which were both expensive to develop and

difficult to maintain. However, recent advances in industry-standard multiprocessor servers that run Windows and SQL Server, such as those available from Dell, have allowed SQL Server running on a Dell™ enterprise server to replace traditional VLDB options.

Rosetta Genomics, a small genomic-data analysis company, is among the first to use standards-based Dell hardware and Microsoft software for a VLDB application. Using the high-performance SQL Server 2000 Enterprise Edition platform, Rosetta Genomics has developed a 10 TB VLDB that has 20 billion records containing human genome information. SQL Server 2000 Enterprise Edition provides advanced features that work well for VLDB implementations. Parallel I/O, automated query parallelism, an in-process BULK INSERT statement, parallel index creation, self-management, and automated memory configuration features provide the required performance for database-intensive applications.

Dell, EMC, Emulex, Microsoft, and Rosetta Genomics collaborated to test the performance effects of several

configuration alternatives for building and loading the Rosetta Genomics™ database. Information gathered through this testing, which was conducted at Microsoft labs, was used to provide setup and configuration parameters to Rosetta. The most effective solutions are presented in this article.

Setting up the test environment

The tests described in this article used a Dell PowerEdge™ 8450 enterprise server running Microsoft SQL Server 2000 Enterprise Edition with Service Pack (SP) 2 on the Microsoft Windows® 2000 Datacenter Server operating system with SP2. The PowerEdge 8450 had eight Intel® Pentium® III Xeon™ processors running at 900 MHz and 32 GB of RAM.

Storage consisted of two EMC® CLARiiON® FC4700 storage arrays and two EMC Symmetrix® 8530 storage arrays. Each of the EMC CLARiiON FC4700 storage arrays had a 4 GB write-read cache, while each EMC Symmetrix 8530 storage array had a 16 GB write-read cache. Each of the four storage arrays included 96 disks capable of 73 GB and 10,000 rotations per minute (RPM), for a total of 28 TB of raw disk space (12 TB after RAID-1 + 0 configuration). The CLARiiON and Symmetrix storage arrays ran EMC PowerPath® (version 2.0 on Symmetrix and version 3.0 on CLARiiON) and EMC ControlCenter™ family 5.0. All of the tests discussed in this article are based on a RAID-1 + 0 configuration of the physical drives in the storage arrays.

Additional hardware included eight LP9002 Emulex® host bus adapters (HBAs). Storage arrays were connected to the PowerEdge 8450 server using two EMC Connectrix™ ED-1032 1 Gbps Fibre Channel switches, and one EMC DS-16B 2 Gbps Fibre Channel switch.

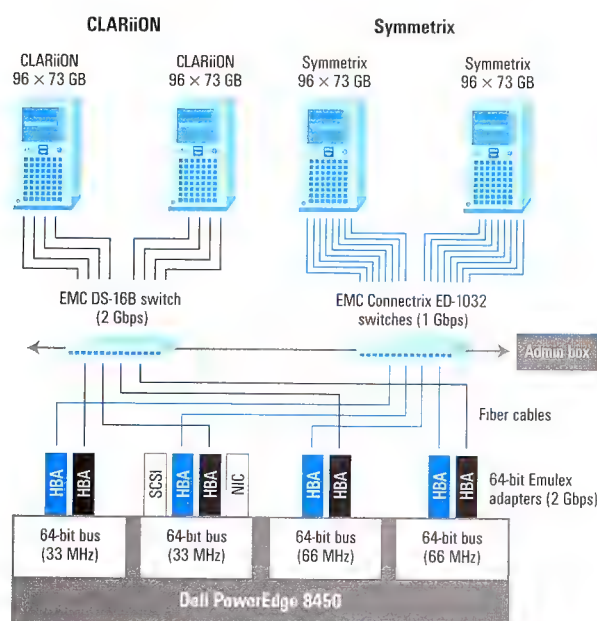


Figure 1. The Rosetta system PCI bus and HBA configuration balances I/O

Setup for the application involved complex preprocessing and analysis of raw genomic sequence data, building the database, and then loading 20 billion records. Rosetta Genomics provided genomic data for the tests. The resulting database was then used as a genomic-sequence data-pattern analysis tool.

Building the human genome database entailed the following general steps:

- Configuring the server Peripheral Component Interconnect (PCI) slots for I/O balancing
- Setting up storage area network (SAN) storage arrays
- Defining Windows 2000 volume mount points
- Creating the database
- Bulk-loading the initial human genome data set
- Extracting proprietary data sets to the file system
- Bulk-loading data sets into the database
- Performing data transformation using the SELECT INTO statement
- Creating the indexes

Configuring server PCI slots

To optimally configure I/O-intensive applications, administrators must balance throughput across as many I/O paths as possible. Achieving even I/O distribution requires an understanding of hardware architecture.

A Dell PowerEdge 8450 server provides four 64-bit PCI buses, and each PCI bus can support up to two 66 MHz PCI cards or up to four 33 MHz PCI cards. To increase I/O potential in this deployment, eight Emulex HBAs were installed, two on each PCI bus. The factory-installed SCSI controller card and network interface card (NIC) were in slots 1 and 7, respectively. By moving the cards to slots 3 and 6, the test team could make available two slots per bus for HBA use, thereby evenly distributing all eight HBAs across the four PCI buses (see Figure 1).

Setting up SAN storage arrays

Spreading a single database across multiple storage arrays increases operational flexibility. Storage for the Rosetta project consisted of 28 TB of raw disk storage across the four storage arrays (two EMC CLARiiON FC4700 arrays and two EMC Symmetrix 8530 arrays). Configuration for redundancy using a RAID-1 + 0 (mirrored and striped disks) format produced 12 TB of available storage. The mixed SAN environment demonstrates how multiple storage units can be added to the SAN, without regard to the storage vendor.

The setup and configuration of the storage arrays' physical drives used logical drives based on as many physical spindles as possible. This approach balanced I/O across spindles and helped eliminate overutilization of any single spindle. Balancing database reads and writes across all drives maximizes I/O bandwidth.

Details of the steps performed to configure the physical drives of the storage arrays will vary depending upon the system design. Because such large databases typically push hardware technology to its limits, Microsoft strongly recommends working closely with hardware vendors when designing terabyte-sized database storage.

Defining Windows 2000 volume mount points

By default, each logical drive on a Windows-based system is presented as a different drive letter. The Windows 2000 operating system introduced volume mount points, which allow a disk volume to be mounted at a particular directory in the structure of another volume without consuming a drive letter.¹

The two CLARiiON and two Symmetrix storage arrays contained 94 logical unit numbers (LUNs). Using volume mount points allowed the test team to use all 94 LUNs presented by the storage array even though only 23 drive letters remained on the system. Administrators ran scripts from the command line using mkdir.exe and diskpart.exe (included in the Windows 2000 Resource Tool Kit) to create the folders, to partition the LUNs, and to assign mount-point directories.

Each mount-point folder was then individually formatted as an NT file system (NTFS) using 64 KB allocation units with format.exe /A:64 for improved SQL Server 2000 disk I/O performance. For large scans, SQL Server 2000 reads data from disk using extents that comprise eight 8 KB data pages; formatting using 64 KB allocation units aligns SQL Server I/O structures with the underlying file system, improving disk I/O potential.

Creating the database

Database creation requires I/O bandwidth across all components of the I/O infrastructure. Database creation has low CPU utilization and high I/O requirements, with most of the overhead taken up in writing the data files. When a database is created, all the files that constitute the database are filled with zeros to overwrite any existing data left on the disk by previously deleted files. Although the files take longer to create, this action prevents the operating system from having to fill the files with zeros as data is written to the files for the first time.

The Rosetta Genomics application database files totaled 10.35 TB for data files and 135 GB for the log file. Administrators defined both a primary, or default, file group and a user-defined file group. The default file group had one data file containing database system objects. Located on the default drive C, the default file group shared drive space with SQL Server binaries and the operating system. The user-defined file group was configured with 40 fixed-size 131 GB data files on CLARiiON arrays and 40 fixed-size 135 GB data files on Symmetrix arrays, for a total of 80 data files—one for each LUN designated for database file storage. Mount points were configured at 132 GB on CLARiiON arrays and 136 GB on Symmetrix arrays.

Duration	Sustained throughput	Average throughput
7 hours, 13 minutes	499 MB/sec	424 MB/sec

Figure 2. Throughput performance metrics for the CREATE DATABASE statement

Through the use of file groups, SQL Server enables administrators to place individual tables and indexes on specific files. Determining which tables and indexes should use separate spindles often helps to isolate and improve access to very active tables or indexes. If, as in the case of the Rosetta Genomics application, administrators cannot predict what type of data access will occur or when, the safer option is to spread the user tables and their indexes across all files in the user-defined file group; doing so helps guarantee consistent performance because all disks are accessed for every operation, while the overhead of maintaining individual file groups is minimized.

After defining file groups, administrators launched CREATE DATABASE as a Transact-SQL statement and included definitions for the 80 data files. The CREATE DATABASE statement is a multi-threaded process that creates one process thread for each logical device, maximizing I/O potential across all of the LUNs. When initializing a database, SQL Server uses a single thread per LUN. By creating one file per LUN and a count of LUNs that provide a whole-number multiple of the available CPUs, the team evenly distributed CPU process threads.

Effect of EMC PowerPath software

EMC PowerPath software balanced I/O for each LUN across multiple HBAs, thus fully utilizing all HBAs during database creation and other I/O-intensive processes. Without PowerPath, a given LUN will use an assigned HBA. LUNs assigned to the 66 MHz HBAs could therefore complete quickly, leaving the 66 MHz HBAs idle while the slower 33 MHz HBAs remained active. Fully employing all available HBAs during I/O-intensive processes improves maximum overall throughput.

CREATE DATABASE test results

During the Rosetta Genomics testing, the SQL Server CREATE DATABASE statement achieved sustained throughput of 499 MB/sec (see Figure 2) and reached 79 percent of the system hardware theoretical maximum capacity of 631 MB/sec.

The test team recommends that for all data and log files, administrators specify an initial size large enough to hold the expected amount of data. Setting initial file size helps administrators avoid the negative performance impact of file growth activity while loading the database.

By building LUNs of the correct size, the team created a desirable number of LUNs and database files without sacrificing

¹ Windows 2000 Cluster Service does not support mount points.

performance. Hardware vendors can assist administrators in determining the optimal LUN size for their applications.

Bulk-loading data from text files

Unlike CREATE DATABASE, each SQL Server 2000 BULK INSERT statement is a single-threaded process, because each statement is directed at a single data file. To achieve parallel data loading, therefore, multiple BULK INSERT statements must be executed, each with a separate source data file. As BULK INSERT streams are started, SQL Server assigns the processes to an available CPU, based on CPU usage, and loads each CPU in turn as data streams are added.

During the Rosetta testing, administrators maintained eight concurrent streams, because the testing included eight CPUs. Keeping the number of bulk insert streams less than or equal to the number of available CPUs improves efficiency, because it reduces the amount of CPU context switching.

When bulk-loading character-based data files, as in the Rosetta project, the data must first be converted to the proper SQL Server data type for a given field before it is loaded into the database. Because of the data conversions, BULK INSERT processes are very CPU intensive, with each stream driving its respective CPU to 90 percent average utilization for this implementation.

In the Rosetta testing, the data-loading process began by using the SQL Server BULK INSERT statement to bulk-load an initial data set from a text file provided by Rosetta Genomics. After the initial load, 16 new data sets were extracted from the initial load and stored on the local file system as text files, using a proprietary Rosetta Genomics process. The 16 text files were then bulk-loaded into the 16 new tables, using the Transact-SQL BULK INSERT statement.

BULK INSERT test results

Figure 3 shows performance metrics for a bulk insert of 16 application-generated flat files using the BULK INSERT statement. In the figure, throughput specifications are provided as an average for the 16 tables.

Rosetta Genomics chose to use the Transact-SQL BULK INSERT statement, because the BULK INSERT statement executes in process with the sqlservr.exe executable, sharing the same memory space. Because the data files are opened by the sqlservr.exe executable, cross-process data copying is avoided.

Duration	BULK INSERT rows/sec	BULK INSERT throughput	Total records
63 minutes	528,009	73.6 MB/sec	1,652,567,235

Figure 3. Throughput performance metrics for the BULK INSERT statement

Duration	SELECT INTO rows/sec	SELECT INTO throughput	Tables created	Total records
57 hours, 12 minutes	914,883	59 MB/sec	1,824	188,392,664,790

Figure 4. SELECT INTO performance metrics

BULK INSERT performance was achieved by following these best practices:

- Set /3GB to Off in the Windows boot.ini file (default).
- Set AWE to Off and Fiber Mode to Off in the SQL Server configuration setting (default).
- Set the SQL Server affinity mask to =0xff (use CPUs 0 through 7). Explicitly setting the affinity mask to use available CPUs reduces context switching, eliminating the cost of level 2 (L2) cache transfers incurred when a process migrates from one CPU to another.
- Keep the number of parallel streams of data equal to the number of CPUs, to optimize CPU utilization while minimizing context switching.
- Implement the BULK INSERT parameter TABLOCK to permit system use of bulk update locks when indexes are absent. This action reduces lock manager overhead and allows the logging of entire data page changes rather than individual records.
- Implement the BULK INSERT parameter CODEPAGE='RAW'. Using CODEPAGE='RAW' makes code page conversion unnecessary and is the fastest option. When using CODEPAGE='RAW' on a bulk copy-in operation, administrators must ensure that, for all character values less than 32 or greater than 127, the collation specifications for SQL Server are the same collations present in the bulk copy data file.
- Use the BULK INSERT default BATCHSIZE parameter and insert data in the entire flat file as a single batch.
- Use the Bulk-Logged Database recovery model, which eliminates record-level and point-in-time logging for bulk insert processes. The Bulk-Logged Database recovery model provides protection against media failure, best performance, and minimal log space usage for certain large-scale or bulk-copy operations.

Performing data transformation using SELECT INTO

The final loading process used data from the 16 tables previously loaded from flat files. After that, the data was selected, transformed, and partitioned using Transact-SQL SELECT INTO statements. Transact-SQL SELECT INTO processing requires both reads and writes to the database and, in the case of the Rosetta Genomics application, data transformation. The transformed and partitioned data was then loaded into the 1,824 new tables required by the Rosetta application. This step increased the data size from 221 GB to 9.97 TB.

SELECT INTO test results

As with BULK INSERT, the SELECT INTO processes can be very CPU intensive. By establishing continuous parallel execution of eight SELECT INTO statements, the testing team achieved the results shown in Figure 4, where each stream drove its respective CPU to 92 percent average utilization.

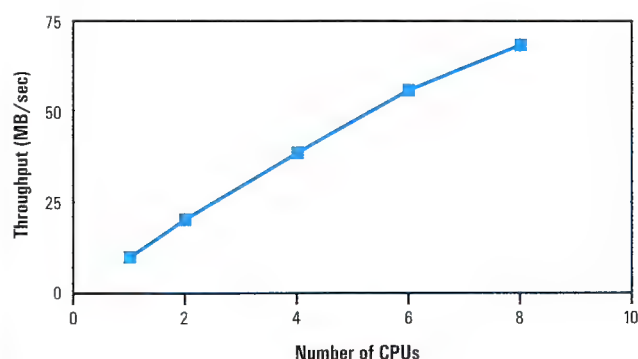


Figure 5. Effect of multiple SELECT INTO streams on throughput

Figure 5 shows scalability testing of SELECT INTO based on the number of CPUs. The figure is based on a test set of data and shows that SQL Server scales almost linearly when adding CPUs and data stream pairs. The team employed the Windows boot.ini NUMPROC parameter to control the number of processors used for this test.

Creating the indexes

Lastly, administrators created 128 nonclustered indexes and 64 clustered indexes on 128 tables. Of these, 64 tables each have a clustered index on two columns as well as a nonclustered index (also on two columns), while the other 64 tables each have only a single-column, nonclustered index. The team chose indexes to support the limited set of queries executed during the workload tests.

CREATE INDEX test results

The number of available CPUs and the configuration option named “max degree of parallelism” determine the maximum number of processors employed to execute a single CREATE INDEX statement. Setting the “max degree of parallelism” option to zero (the default) allows SQL Server to dynamically determine the optimal degree of parallelism up to the maximum number of available CPUs. If SQL Server detects that the system is busy, the degree of parallelism of the CREATE INDEX operation is automatically reduced before statement execution begins. In very large data loading scenarios such as this one, administrators typically perform index creation while the server is otherwise idle, allowing the CREATE INDEX process to take advantage of greater CPU utilization (see Figure 6).


The CREATE INDEX performance metrics reflect the ability of SQL Server 2000 Enterprise Edition to create indexes in parallel across multiple processors.

Duration	Processor time	Peak disk read bytes/sec	Peak disk write bytes/sec	Data indexed	Indexed size
15 hours, 20 minutes	82%	97 MB/sec	85 MB/sec	350 GB	337.79 GB

Figure 6. CREATE INDEX performance metrics

A scalable, low-cost solution for VLDBs

Microsoft SQL Server 2000 Enterprise Edition running on the Windows 2000 Datacenter Server operating system provides excellent performance for very large databases using readily available, standard hardware, such as Dell servers. SQL Server 2000 and Windows 2000 Datacenter Server achieved outstanding performance during database creation and data loading using default settings and dynamic configuration capabilities. Transact-SQL BULK INSERT and SELECT INTO statements offer exceptional performance through implementation of parallel data streams and clearly demonstrate the ability of SQL Server 2000 Enterprise Edition to scale up.

The Dell and Microsoft solution used by Rosetta Genomics costs a fraction of comparable supercomputing solutions and provides affordable scalability as the genomic database grows. Both initial cost structure and ongoing operational costs are lower than those of mainframe and supercomputer implementations, helping IT organizations that support VLDB applications to lower their TCO. 

Acknowledgments

Content for this article was the result of joint efforts of the following companies: Dell, EMC, Emulex, Microsoft, and Rosetta Genomics. Dell, EMC, and Emulex provided both the equipment and technical resources for the installation and configuration of their respective hardware, which aided in the successful deployment and performance testing of the 10 TB database described in this article.

Joe Belson (joebel@microsoft.com) is a field consultant for Microsoft Consulting Services (MCS). He specializes in enterprise-level SQL Server and Analysis Server deployments with an emphasis on performance tuning.

Hao Chen (haoc@microsoft.com) is a technical lead in the Microsoft SQL Server Release Services Group, focusing on SQL Server scalability testing. She has a master's degree in Electrical Engineering from Gonzaga University.

FOR MORE INFORMATION

SQL Server:

<http://www.microsoft.com/sql>

Rosetta Genomics:

<http://www.rosettagenomics.com>

EMC enterprise storage platforms:

http://www.emc.com/products/physical_performance.htm

Dell servers:

http://www.dell.com/us/en/biz/products/line_servers.htm

SQL Server on Dell:

<http://www.dell.com/SQL>

Emulex Fibre Channel products:

<http://www.emulex.com/products>

Sizing Microsoft Exchange 2000

on Dell PowerEdge Servers and Dell | EMC Fibre Channel Storage Systems



When deploying and sizing hardware for a Microsoft® Exchange 2000 messaging infrastructure, IT administrators must consider several factors. This article explains how to size Intel® Xeon™ processor-based Dell™ PowerEdge™ servers and Dell | EMC storage hardware for a scalable Exchange 2000 deployment, and then how to tune the operating system and Exchange 2000 Server application.

BY GUY WESTBROOK

The Microsoft® Exchange 2000 Server application can help enterprises build a scalable, secure messaging infrastructure. Before implementing an Exchange 2000 messaging infrastructure, IT administrators should be familiar with the installation and basic administration of Exchange 2000 and the Microsoft Windows® 2000 Advanced Server operating system. The next steps in an Exchange 2000 deployment are to size the server and storage hardware and then to tune the operating system and Exchange 2000 Server application. This article discusses factors contributing to hardware sizing and software tuning, and provides recommendations for those deploying Microsoft Exchange 2000 Server in environments using Intel® Xeon™ processor-based Dell™ PowerEdge™ servers and Dell | EMC Fibre Channel storage systems, such as the Dell | EMC FC4700.

Analyzing Exchange 2000 performance

Several tools can help administrators to size the hardware needed for a messaging infrastructure. Benchmarking tools can assist in comparative analysis of hardware,

the selection of hardware, and the sizing of the hardware components. One of the most popular tools for comparing hardware running Exchange 2000 Server is a tool developed by Microsoft called LoadSim. The new LoadSim MAPI (Messaging Application Programming Interface) Messaging Benchmark 2 (MMB2) measures throughput in terms of a specific profile of user actions, executed over an eight-hour workday.

LoadSim results should be interpreted as a benchmark for messaging throughput and should not be confused with deployment recommendations. Factors such as server and network topology, backup and restore requirements, and other issues should be considered when planning a deployment. All of these factors can have a significant effect on hardware decisions and system performance.

Evaluating the LoadSim test results

In tests conducted at the Dell System Performance and Analysis Lab, the LoadSim tool was used to simulate 11,304 Exchange (MAPI) e-mail users running on a Dell

PowerEdge 6650 server. During this test, the PowerEdge 6650 provided a weighted 95th percentile response time of 137 milliseconds for 11,304 users, with overall CPU utilization of 67 percent. Results were based on four hours of steady-state running.

The PowerEdge 6650 supports large on-die cache—providing up to 1.25 MB per processor—for large-footprint applications such as online transaction processing (OLTP) databases and data warehouses. For memory-intensive applications, the PowerEdge 6650 offers up to 16 GB of SDRAM.

Building an Exchange 2000 architecture

In efforts to determine best practices for the Exchange 2000 Server application, the Dell testing team used a Dell PowerEdge 6650 as the Exchange 2000 server. Eight Dell PowerEdge 6450 servers were used as load simulator clients. The team configured the PowerEdge 6650 with four Intel Xeon processors MP at 2.0 GHz and with 2 MB of L3 cache and 4 GB of memory.

The team also used one Dell | EMC FC4700 storage array and two Dell | EMC disk array enclosures (DAEs) as the Fibre Channel storage subsystem for the Exchange database. All enclosures were populated with 36 GB 10,000 rpm Fibre Channel disks, and the FC4700 included two storage processors, each with 1 GB of on-board configurable cache and two Intel Pentium® III processors at 733 MHz. For storage subsystem connectivity, the Exchange server contained two Emulex® 9002 host bus adapters (HBAs). Two Dell | EMC DS-16B2 Fibre Channel 2 switches were used for the switched fabric.

Figure 1 depicts this test configuration, which was used to gather data for this article. The configuration included fully redundant paths for high availability.

Understanding I/O factors affecting an Exchange deployment

The disk I/O subsystem is the most critical component in sizing an Exchange environment. The goal of performance-tuning the Exchange environment is to decrease response time while supporting more users. This tuning can be accomplished by adding more physical disk drives or spindles.

Exchange 2000 Server uses random I/O and 4 KB to 7 KB transfer sizes for the database. Transaction log files are 100 percent sequential I/O and use 4 KB transfer writes. Write efficiency improves through using a write-back cache. The Dell | EMC storage subsystem can provide on-board cache of up to 1 GB per storage processor. The cache is write-mirror protected and both read and write cache can be used for the read I/O. Because the I/O patterns for the transaction log and the database are different, they should be configured and tuned differently.

Several tools can help administrators to size the hardware needed for a messaging infrastructure.

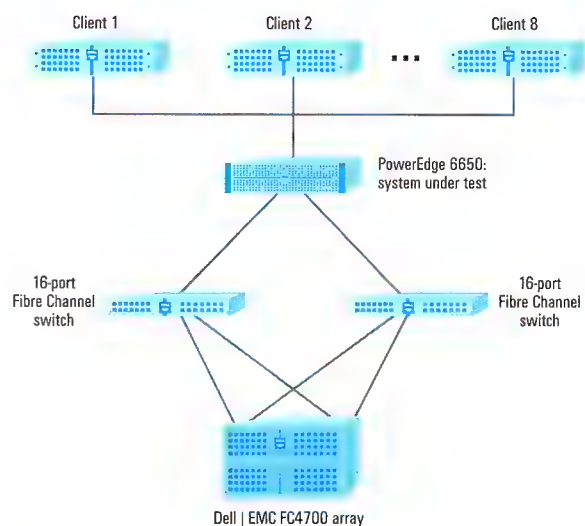


Figure 1. Microsoft Exchange messaging infrastructure tested at Dell System Performance and Analysis Lab

Optimizing the transaction log

Transaction logs are important because they are the key to recovering an Exchange database. For a small environment such as one with only a few hundred users, the effect of transaction log optimization is not significant. But for larger environments, optimizing transaction log files can provide performance gains and prevent bottleneck issues relating to the transaction logs. Transaction log bottlenecks can degrade performance of the entire Exchange environment. Recommended best practices for optimizing transaction logs include:

- Using a dedicated logical unit number (LUN) for the transaction log files
- Avoiding multiple transaction log files on the same LUN
- Enabling write-back mode on the LUN
- Using RAID-1 (two spindles) or RAID-5 (three spindles) for transaction logs

One optimization parameter that can be tuned to increase transaction log performance is the transaction log buffer size. This tuning is performed from the server side, not at the cache on the back-end storage system. Exchange 2000 includes a default of 84 buffers for transaction log files. The database engine uses these buffers to store transactions until they can be committed to the disk. In a large environment in which I/O might be busy, the buffers can fill up, causing write delays.

Transaction log buffers are storage group-based. To optimize the log buffer, use ADS/Edit, the Microsoft Active Directory® Service Interfaces (ADSI) Management Console snap-in, and find the storage group. Then find the msEXchESEParamLogBuffers object



The goal of performance-

tuning the Exchange

environment is to

decrease response

time while supporting

more users.

and change the size from 84 to 9,000 buffers.

Optimizing the database file

The I/O pattern for the Exchange database—random reads and writes with an approximate 60/40 split—is different from that of the transaction log. Based on this I/O pattern, the read cache does not significantly affect database-access performance, but using the write cache will decrease response time, improving performance. Administrators can change the memory distribution of each storage processor in the Dell | EMC FC4700 based on this I/O pattern. The FC4700 has 1 GB of cached memory: 50 MB should be set for read caching and the rest for write caching, because I/O requests use both read cache and write cache to fetch data from disk.

To increase the number of concurrent supported users in an Exchange environment, the spindle count should be increased. In addition, other factors influence sizing the Exchange database. Administrators must determine how many users will be supported and how large user mailboxes can grow. The deleted items retention policy also requires storage space, which should be approximately 25 percent additional space. Finally, whether circular logging is turned on or off should be factored into the total space requirements.

Establishing RAID level and stripe size

RAID not only provides data protection, but it also can improve performance by adding more drive spindles to the array group. Adding drives to an array provides greater throughput regardless of storage capacity and fault-tolerance level. Environments with more random and small-block transfers, such as messaging, will enjoy greater performance benefits with more drive spindles.

Because Exchange 2000 Server is a write-intensive application, the recommended RAID configuration on the drives containing the mail database files is RAID-5. However, RAID configurations may differ based on unique customer environments and situations. If performance is the only factor to be considered and fault tolerance is not required, RAID-0 will best meet the needs of the environment. RAID-0 is well suited for tasks such as disk backup, because RAID-0 can provide faster I/Os than other RAID levels and fault tolerance is not required. However, most messaging application deployments demand fault tolerance for mail databases and transaction logs. In these cases, RAID-5 provides an effective combination of performance, economy, and reliability.

Setting up the I/O subsystem for Exchange 2000

In the sizing tests, the Dell team used RAID-5. Although documentation recommends using the default stripe size, which is 64 KB

(128 elements), tests show that a 32 KB stripe size performs slightly better than the default 64 KB.

Figure 2 shows test results for five drives running RAID-5 in a 4+1 RAID set. At this number of physical drives or spindles, the maximum recommended number of MMB2 users for maintaining a subsecond response time was 2,000. When the number of MMB2 users exceeded 2,000, the response time increased to more than 4 seconds. Figure 2 also shows test results for nine drives running RAID-5 in an 8+1 RAID set. To achieve a subsecond response time, the maximum number of users recommended for this number of drives was 4,000.

For implementations that require more than 4,000 MMB2 users, the LUN design should be planned so that it corresponds with usage requirements. The following formula can be used to determine the number of drives required for a specific number of users:

$$\begin{aligned} 150 \text{ I/Os per drive} \times \text{Number of drives} &= \text{Total I/Os per second} \\ \text{Total I/Os per second} / 0.35 &= \text{Number of users} \end{aligned}$$

This formula assumes that the recommended number of I/Os per disk drive is 150, because more than 150 I/Os can cause a significant increase in average response time (ART) and can result in unstable system behavior. The recommended number of I/Os for each user (0.35) was derived from MMB2 testing. *Number of users* is the total number of users that can be supported by the RAID level and the number of drives while still achieving a subsecond response time.

Sizing the hardware for an Exchange 2000 deployment

When sizing hardware, server capabilities should be considered. The following factors will affect hardware requirements:

- Number of users and their usage patterns
- Concurrent versus nonconcurrent users (the typical concurrency

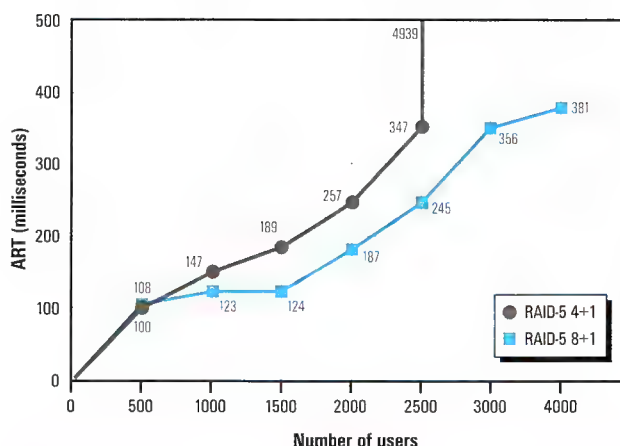


Figure 2. Average response time (ART) with five-drive and nine-drive RAID-5 arrays

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(0)\WINNT
[operating systems]
multi(0)disk(0)rdisk(0)partition(0)\WINNT="Windows 2000 Advanced Server [3GB user]" /3GB
4.00 [VGA mode]" /basevideo /sos
```

Figure 3. Modifying the boot.ini file to enable the 3 GB user memory feature

rate for messaging servers is 60 to 65 percent of the total number of users)

- Processor performance and availability of Hyper-Threading technology (Hyper-Threading can provide a 20 to 25 percent boost in performance, according to Dell tests)
- L2 and L3 cache sizes
- Maximum amount of RAM
- Number of Peripheral Component Interconnect (PCI) slots
- Number and speed of network interface cards (NICs)
- Number of drive controllers
- Number, capacity, and speed of hard drives
- Stripe size for data (the recommended size is 32 KB)
- RAID configuration and fault-tolerance requirements
- Availability, scalability, and failover requirements (clustering)
- Backup and restore requirements
- Response-time criteria

Identifying important server components

Most performance problems occur because of insufficient hardware, improper hardware configuration, or an unoptimized Exchange Server configuration. Sizing the system usually involves the CPU, memory, I/O subsystem, and other related hardware components.

PCI bus. The PCI bus is a high-performance, 32-bit or 64-bit local bus that provides an interface for high-speed transfers to peripheral components without overloading the host processor. When installing an even number of network or array controllers, administrators should split the controllers evenly between the buses. In general, network controllers consume more bandwidth than array controllers.

Processor scalability. The processor subsystem allows the server to handle an increasing number of users without compromising response time. Therefore, administrators should use the fastest processor with the largest amount of cache. Such processors will support better initial performance and future growth.

Tuning the operating system and the Exchange 2000 application

By default, Microsoft Windows 2000 Advanced Server permits a maximum 2 GB of the installed memory to be allocated as user memory. However, after installing Microsoft Windows 2000 Advanced Server, administrators can modify the system so that in a server with 4 GB memory, up to 3 GB of memory can be consumed by user applications such as Exchange 2000. The remaining 1 GB of a 4 GB system would then be reserved for the Windows 2000 Advanced Server system and unavailable as user memory. (Windows 2000 can use more than 4 GB of RAM with Physical Address Extension [PAE] support, but currently Exchange 2000 Server does not support PAE.)

To enable the 3 GB user memory feature, administrators must edit the hidden system file c:\boot.ini and add the command-line option /3GB to the boot-up string (see Figure 3), and then reboot the system.

If a server is supporting a large number of users, administrators can specify that Windows 2000 needs approximately 64 MB for itself. Included in this 64 MB will be some memory allocated for the system cache. Administrators should then minimize the system cache size by selecting "Maximum Data Throughput for Network Applications." This setting will cause Windows 2000 to page memory used by the system cache before it pages memory used by Exchange 2000 Server. If the system is running low on

memory resources, the Exchange 2000 Server processes will begin to have their memory paged. This situation is usually the first warning sign that a server needs more memory.

RAID not only provides data protection, but it also can improve performance by adding more drive spindles to the array group.

Factors affecting Exchange 2000 memory usage

The recommendations for Exchange 2000 RAM requirements vary between 300 KB to 1 MB per user, depending on the user's



activities. During benchmarking, Dell engineers have used the following formula for determining memory requirements:

$$256 + (\text{Number of concurrent users} / 5) = \text{Required RAM (MB)}$$

Usage patterns and whether the users work directly on the server rather than replicating databases to their own computer (and thus freeing processing resources on the server) will also affect the amount of required memory.

A key optimization parameter that can be set is the Mailstore Database Cache Size on Exchange 2000, which has a default value of 900 MB. On servers with more than 2 GB of memory, increasing the size of this cache may be beneficial. Because of virtual address space limitations, this value should never be set higher than 1200 MB.

The 900 MB default limit helps ensure that the store.exe process always has ample virtual address space (memory) from which to allocate. Increasing this value too much can lead to system instability. The information gathered from the Performance Monitor counter, Virtual Bytes, will provide an accurate value for the virtual address space that the store.exe process has allocated. On a server with the /3GB switch set in the boot.ini file, the value for Virtual Bytes should be below 2.8 GB. On a server without the /3GB switch set, the value for Virtual Bytes should be below 1.8 GB.

If values for either configuration are higher, administrators should not increase the maximum cache size. If values for either configuration are lower, administrators can safely increase the size of the database maximum cache size. To modify the cache size, administrators should use the Microsoft ADSI tool, select the msExchESEParamCacheSizeMax attribute, and adjust the value.

Transaction logs are
important because
they are the key
to recovering an
Exchange database.

Creating a robust messaging infrastructure

By appropriately sizing the back-end storage and by optimizing the operating system and Exchange application, administrators can combine Dell PowerEdge servers, Dell | EMC Fibre Channel storage systems, and Microsoft Exchange 2000 Server to create a robust messaging infrastructure in an enterprise environment. Exchange 2000 Server takes full advantage of the high availability and scalability delivered by such a configuration. ➤

References

EMC. "EMC CLARiiON Data Integrity Difference."

http://www.emc.com/pdf/products/clarion/clarion_data_integ.pdf.

Hou, Richard. "Running an Optimized Microsoft Exchange 2000 Server with Dell | EMC SAN." *Dell Power Solutions*, May 2002.

Microsoft Corporation. *Microsoft Exchange 2000 Server Resource Kit*. Redmond, Wash.: Microsoft Press, September 2000.

Microsoft Corporation. "Microsoft Exchange Server: Performance and Scalability." <http://www.microsoft.com/exchange/techinfo/planning/2000/PerfScal.asp>.

Stanford, Scott. "The Effect of L3 Cache Size on MMB2 Workloads." *Dell Power Solutions*, February 2003.

Guy Westbrook (guy_westbrook@dell.com) is a performance engineer in the System Performance and Analysis Lab at Dell. He is responsible for performance benchmark testing of enterprise messaging applications, including Lotus® Domino™ and Microsoft Exchange. Guy has a B.S. in Computer Science from the University of Arkansas at Pine Bluff and an M.S. in Information Systems from Western International University. He is also a Microsoft Certified Systems Engineer + Internet (MCSE+Internet).

FOR MORE INFORMATION

Microsoft Exchange: <http://www.microsoft.com/exchange>

Dell | EMC storage: <http://www.dell.com/emc>

Share Your
Experience in
Dell Power Solutions

Dell Power Solutions is a peer-to-peer communication forum. We welcome subject-matter experts, end users, business partners, Dell™ engineers, and customers to share best-practices information. Our goal is to build a repository of solution white papers to improve the quality of IT.

Guidelines for submitting articles to *Dell Power Solutions* can be found at <http://www.dell.com/powersolutions>

Microsoft and Dell

Out in front



Microsoft® SQL Server™ has
transformed the database management landscape
in the 10 years since its inception.

By continually introducing innovative new features, SQL Server has set the bar ever higher for relational database performance—while keeping the total cost of ownership low.¹ Independent benchmarking proves that Dell™ PowerEdge™ servers can contribute to this leading low-cost, high-performance solution.²

To find out more, visit us online at www.microsoft.com/sql and www.dell.com/sql.



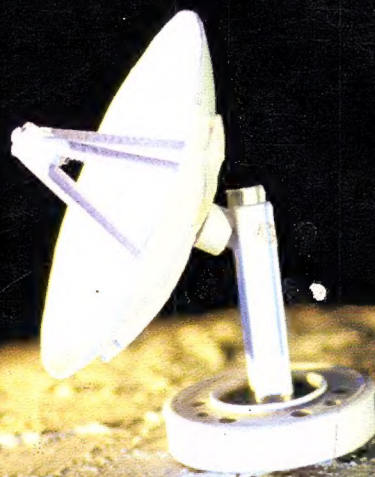
Microsoft
SQL Server 2000
Enterprise Edition

¹For more information, visit <http://www.microsoft.com/sql/evaluation/anniversary/roadmap.asp>

²Dell is a consistent leader in the price/performance category of TPC-C benchmarks running Microsoft SQL Server 2000 on Dell PowerEdge Servers. To learn more about TPC-C benchmarking, visit the Transaction Processing Performance Council Web site at www.tpc.org.

© 2003 Microsoft Corporation and Dell Computer Corporation. All rights reserved. Microsoft is a registered trademark of Microsoft Corporation. Dell, the Dell logo, and PowerEdge are trademarks of Dell Computer Corporation. Other company, product, and service names may be trademarks or service marks of others.

Who says **SIZE** doesn't matter?



When your IT organization has
terabytes of data to process,

YOU'VE GOT TO THINK BIG.

Microsoft and Dell understand this.

A proven combination, Microsoft® SQL Server™ running on Dell™ PowerEdge™ servers can handle billions of records with impressive performance. More companies than ever entrust their biggest jobs to this powerful enterprise solution—making SQL Server on Dell hardware a fast-growing solution for very large databases.

Of course bigger isn't always better. That's why this major processing power is available for a minimal total cost of ownership. To learn more, visit www.microsoft.com/sql and www.dell.com/sql.

© 2003 Microsoft Corporation and Dell Computer Corporation. All rights reserved. Microsoft is a registered trademark of Microsoft Corporation. Dell, the Dell logo, and PowerEdge are trademarks of Dell Computer Corporation. Other company, product, and service names may be trademarks or service marks of others.

Microsoft
SQL Server 2000
Enterprise Edition